CHAPTER 15

NETWORK SECURITY

After reading this chapter and completing the exercises, you will be able to:

- Identify security risks in LANs and WANs
- Explain how physical security contributes to network security
- Discuss hardware- and design-based security techniques
- Use network operating system techniques to provide basic security
- Implement enhanced security through specialized software
- Describe the elements of an effective security policy

ON THE JOB

With few exceptions, today's commercial computer networks are vulnerable to break-ins. As a network security assessor, I've seen even the most secure commercial networks breached within a few days' time.

Commercial state-of-the-art network security includes such technical solutions as public key encryption, one-time passwords, over 1024-bit strong encryption, fingerprint verification, and even retinal scanning. But the reality remains that an intruder can access secured systems by finding "convenient short cuts" installed by system and network administrators.

During network security assessments, for example, we often find a trusted relationship between two or more hosts. One of the hosts is usually less secure than the others. When we gain access to one system, the others dutifully allow us right in.

In one instance, a firewall-protected network appeared very secure. However, when we scanned the TCP port ranges, we found the signature of a popular commercial mail server. We ran a mail client, attached to the server, and tried a couple of generic user IDs. Incredibly, we were able to gain access to the network manager's mailbox without a password! Even worse, sensitive mail was not encrypted. One message included the complete set of configuration statements for some remote office network routers. Without firewalls, these routers were connected to the Internet and had trusted relationships with the home office. On gaining access to one of the remote office routers, we immediately logged in to the organization's home office router. From there, it was a short time before we found and gained access to a number of UNIX and Windows NT servers—the backbone of the organization's network.

Two small "conveniences" enabled us to circumvent the highly secured firewall and breach the organization's network within a couple of hours!

David Klann
Berbee Information Networks, Inc.

In the early days of computing, when secured mainframes acted as central hosts and data repositories that were accessed only by dumb terminals with limited rights, network security was all but unassailable. As networks have become more geographically distributed and heterogeneous, however, the risk of their misuse has also increased. Consider the largest, most heterogeneous network in existence: the Internet. Because it contains millions of points of entry, millions of servers, and millions of miles of cabling, it is vulnerable to millions of break-ins. Because so many networks connect to the Internet, the threat of an outsider accessing an organization's network via the Internet, and then stealing or destroying data, is very real. In this chapter, you will learn how to assess your network's risks, how to manage those risks, and, perhaps most importantly, how to convey the importance of network security to the rest of your organization through an effective security policy.

TERMINOLOGY

Before delving into network security issues, you should have a clear understanding of terminology frequently used in this field. First, you should understand the difference between a hacker and a cracker. A **hacker** is someone who masters the inner workings of operating systems and utilities in an effort to better understand them. A **cracker** is someone who uses his or her knowledge of operating systems and utilities to intentionally damage or destroy data or systems. The primary difference between hackers and crackers is that hackers do not conduct their experimentation with malicious intent. In fact, hackers may be commissioned to break into networks as part of security audits, in an effort to test whether a cracker could do the same. This chapter will focus on network security in terms of how it protects against crackers.

Another frequently used term related to network security is root. In general, **root** refers to a highly privileged user ID that has all rights to create, delete, modify, move, read, write, or execute files on a system. Specifically, it may mean the administrator on a UNIX-based network. Getting the root ID and password on one system often allows crackers to gain access to attached systems, which is typically their goal. For this reason, information about root accounts should be carefully guarded.

Every network operating system requires that users provide authentication. As you learned in Chapter 8, authentication is the process of verifying a user's validity and authority on a system; it generally takes place during the login process. Different systems use different credentials to authenticate users as they log on. You are probably most familiar with the user ID and password combination. Some systems, however, may also base authentication on digital signatures, IP addresses, session IDs, or a combination of these methods. Generally, the more information required for authentication, the stronger the authentication, and the more secure the system. If a protocol or system uses little information to verify an attempt to access its data, the protocol or system is considered to have weak authentication.

One more term you should understand is "firewall". A **firewall** is a specialized device (usually a router, but possibly only a PC running special software) that selectively filters or blocks traffic between networks. A firewall typically involves a combination of hardware and software (for example, the router's operating system and configuration). The term "firewall" is derived from the physical "wall" installed between rooms in a building or in automobiles between the passenger area and the engine to help prevent fires from spreading from one space to another.

SECURITY AUDITS

Before spending time and money on network security, you should examine your network's security risks. As you learn about each risk facing your network, you should consider the effect that a loss of data, programs, or access would have on your network. The

more serious the potential consequences, the more attention you will want to pay to the security of your network. In general, deciding how much to invest in network security is similar to buying life insurance. If you are unmarried with no dependents, no debts, and very few assets, you might not care much about life insurance coverage; you might therefore choose not to invest in monthly premiums. If you have six children and you run a business that employs 200 people, however, you will want to ensure that your loved ones and your business will be protected in case of your death.

In much the same way, different types of organizations have variable levels of network security risk. For example, if you work for a large savings and loan institution that allows its clients to view their current loan status online, you must consider a number of risks associated with data and access. If someone obtained unauthorized access to your network, all of your customers' personal financial data would be vulnerable. On the other hand, if you work for a local greenhouse that is not connected to the Internet and uses its internal LAN only to track plant inventory and sales, you may not be concerned if someone gains access to your network, because you have little to lose and nothing is very confidential. Just as you learned in Chapter 14, the key question is, "What will I lose if my system goes down?" In addition, when considering security risks, you should ask, "How much of the information that I store, transmit, and receive is confidential?"

Every organization should assess its security risks by conducting a **security audit** (a thorough examination of each aspect of the network to determine how it might be compromised). You should perform regular security audits at least annually and preferably quarterly; in addition, you should conduct a security audit after making any significant changes to your network. For each threat listed in the following sections, your security audit should rate the severity of its potential effects, as well as its likelihood. A threat's consequences may be severe, potentially resulting in a network outage or the dispersal of top-secret information, or it may be mild, potentially resulting in a lack of access for one user or the dispersal of a relatively insignificant piece of corporate data. The more devastating a threat's effects and the more likely it is to happen, the more rigorously your security measures should address it. Appendix D, "Examples of Standard Networking Forms," provides an example of a checklist you can use to perform a fundamental security audit.

A qualified consulting company can also conduct security audits for your network. The advantage of having an objective third party, such as a consultant, analyze your network is that he or she might find risks that you overlooked because of your familiarity with your environment. Third-party audits may seem expensive, but if your network hosts confidential and critical data, they will be well worth their cost.

After identifying your network's vulnerabilities, you should examine how each security risk might detrimentally affect your data and systems. In the next section, you will learn about risks associated with people, hardware, software, and Internet access.

SECURITY RISKS

Now that you understand the basic terms associated with network security, you are ready to learn about the types of risks facing most networks. The following sections describe these risks. Later in this chapter, you will learn how to protect against each type of threat.

As you learned in Chapter 14, natural disasters, viruses, and power faults can damage a network's data, programs, and hardware. A security breach, however, can harm a network just as easily and quickly. To understand how to manage network security, you should first recognize the types of threats that your network may suffer. Not all security breaches result from a manipulation of network technology. Instead, some occur when staff members purposely or inadvertently reveal their passwords; others result from undeveloped security policies.

As you read about each security threat, think about how it could be prevented, whether it applies to your network (and if so, how damaging it might be), and how it relates to other security threats. Keep in mind that malicious and determined intruders may use one technique which then allows them to use a second technique, which then allows them to use a third technique, and so on. For example, a cracker might discover a user's ID by watching her log onto the network; the cracker might then use a password-cracking program to access the network, where he might plant a program to generate an extraordinary volume of traffic that essentially disables the network's connectivity devices.

Risks Associated with People

By some estimates, human errors, ignorance, and omissions cause more than half of all security breaches sustained by networks. One of the most common methods by which an intruder gains access to a network is to simply ask a user for his or her password. For example, the intruder might pose as a technical support analyst who needs to know the password to troubleshoot a problem, or the password might be learned through a casual conversation about passwords. This strategy is commonly called **social engineering**, because it involves manipulating social relationships to gain access. This and other risks associated with people are listed below. Many people-related risks can be addressed through a clear, simple, and strictly enforced enterprise-wide security policy. You will learn how to develop an effective security policy later in this chapter.

Risks associated with people include the following:

- Intruders or attackers using social engineering or snooping to obtain user passwords
- An administrator incorrectly creating or configuring user IDs, groups, and their associated rights on a file server, resulting in file and login access vulnerabilities
- Network administrators overlooking security flaws in topology or hardware configuration

- Network administrators overlooking security flaws in the operating system or application configuration
- Lack of proper documentation and communication of security policies, leading to deliberate or inadvertent misuse of files or network access
- Dishonest or disgruntled employees abusing their file and access rights
- An unused computer or terminal being left logged onto the network, thereby providing an entry point for an intruder
- Users or administrators choosing easy-to-guess passwords
- Authorized staff leaving computer room doors open or unlocked, allowing unauthorized individuals to enter
- Staff discarding disks or backup tapes in public waste containers
- Administrators neglecting to remove access and file rights for employees who have left the organization
- Users writing their passwords on paper, then placing the paper in an easily accessible place (for example, taping it to their monitor or keyboard)

Human errors account for so many security breaches because taking advantage of them is the easiest way to circumvent network security. Imagine a man named Kyle, who was recently fired from his job at a local bank. Because Kyle felt he was unfairly treated, he wants to take revenge on his employer. He still has a few friends at the bank. Even though the bank's network administrator was wise enough to deactivate Kyle's network logon ID and rights upon his termination, and even though the bank has a policy prohibiting employees from sharing their passwords, Kyle knows his friends' IDs and passwords. Nevertheless, the bank's policy prevents former employees from walking into its offices.

How might Kyle attain his goal of deleting a month's worth of client account activity statements? Although the bank has a network security policy, employees such as Kyle's friends probably don't pay much attention to it. Kyle could most likely walk into the bank's offices, ostensibly to meet one of his friends for lunch. While in the offices, Kyle could either sit down at a machine where his friend was still logged on or log on as his friend because he knows his friend's password. Once in the system, he could locate the account activity statements and delete them. Although this example may be an oversimplification of the process, it isn't far from reality.

Risks Associated with Hardware and Network Design

This section describes security risks inherent in (roughly) Layers 1 and 2 of the OSI Model—the Physical and Data Link layers. Recall that the transmission media, NICs, hubs, and network transmission methods (for example, Ethernet) reside at these layers. This section will also discuss security risks in higher-level hardware, such as routers. At these levels, security breaches require more technical sophistication than those that take advantage of human errors. For instance, to eavesdrop on transmissions passing over

CAT5 cabling, an intruder must use a device such as a sniffer. In the middle layers of the OSI Model, it becomes somewhat difficult to distinguish between hardware and software techniques. For example, because a router acts to connect one type of network to another, an intruder might take advantage of the router's security flaws by sending a flood of TCP/IP transmissions to the box, thereby disabling it from carrying legitimate traffic. You will learn about software-related risks in the following section.

The following risks are inherent in network hardware and design:

- Wireless and wire-based transmissions can often be intercepted (whereas spread-spectrum wireless and fiber-based transmissions cannot).
- Networks that use leased public lines, such as T1s or ISDN connections to the Internet, are vulnerable to eavesdropping.
- Network hubs broadcast traffic over the entire segment, thus making transmissions more widely vulnerable to sniffing. (By contrast, switches provide logical point-to-point communications, which limit the availability of data transmissions to the sending and receiving nodes.)
- Unused hub, router, or server ports can be exploited and accessed by crackers if they are not disabled. A router's configuration port, accessible by Telnet, may not be adequately secured.
- If routers are not properly configured to mask internal subnets, users on outside networks (such as the Internet) can read the private addresses.
- Modems attached to network devices may be configured to accept incoming calls, thus opening security holes if they are not properly protected.
- Dial-in access servers used by telecommuting or remote staff may not be carefully secured and monitored.
- Computers hosting very sensitive data may coexist on the same subnet with computers open to the general public.

While security breaches occur less frequently at the lower layers of the OSI Model, they can prove equally, if not more, damaging than security breaches at the higher layers of the OSI Model. Imagine that a cracker wants to bring a library's database and mail servers to a halt. Suppose also that the library's database is public and can be searched by anyone on the Web. The cracker might begin by scanning ports on the database server to determine which have no protection. If she found an open port on the database server, the cracker might connect to the system and deposit a program that would, a few days later, damage operating system files. Or she may launch a heavy stream of traffic that overwhelms the machine and prevents it from functioning. She might also use her newly discovered access to determine the root password on the system, gain access to other systems, and launch a similar attack on the library's mail server, which is attached to the database server. In this way, even a single mistake (not protecting an open port) on one server can lead to failures of multiple systems.

Risks Associated with Protocols and Software

Like hardware, networked software is only as secure as you configure it to be. This section describes risks inherent in the higher layers of the OSI Model, such as the Transport, Session, Presentation, and Application layers. As noted earlier, the distinctions between hardware and software risks are somewhat blurry because protocols and hardware operate in tandem. For example, if a router has not been properly configured, a cracker may exploit the openness of TCP/IP to gain access to a network. Network operating systems and application software present different risks. In most cases, their security is compromised by a poor understanding of file access rights or simple negligence in configuring the software. Remember—even the best encryption, computer room door locks, security policies, and password rules make no difference if you grant the wrong users access to critical data and programs.

The following are some risks pertaining to networking protocols and software:

- TCP/IP contains several security flaws. For example, IP addresses can be falsified easily, checksums can be thwarted, UDP requires no authentication, and TCP requires only weak authentication.
- Trust relationships between one server and another may allow a cracker to access the entire network because of a single flaw.
- Network operating system software typically contains "backdoors" or security flaws. Unless the network administrator performs regular updates, a cracker may exploit these flaws.
- If the network operating system allows server operators to exit to a command prompt, intruders could run destructive command-line programs.
- Administrators might accept the default security options after installing an operating system or application. Often, defaults are not optimal. For example, the default user ID that enables someone to modify anything in Windows 2000 Server is called "Administrator." This default is well known, so if you leave the default ID as "Administrator," you have given a cracker half the information he or she needs to access your system and obtain full rights.
- Transactions that take place between applications, such as databases and Web-based forms, may be open to interception.

To understand the risks that arise when an administrator accepts the default settings associated with a software program, consider the following scenario. Imagine that you have invited a large group of computer science students to tour your IT department. While you're in the computer room talking about subnetting, a bored student standing next to a Windows 2000 Professional workstation that is logged onto the network decides to find out which programs are installed on the workstation. He discovers that this workstation has the SQL Server administrator software installed. Your organization uses a SQL database to hold all of your employees' salaries, addresses, and other confidential information. The student knows a little about SQL, including the facts that the default administrator

user ID is called "sa," and that, by default, no password is created for this ID when someone installs SQL Server. He tries connecting to your SQL database with the "sa" user ID and no password. Because you accepted the defaults for the program during its installation, within seconds the student is able to gain access to your employees' information. He could then change, delete, or steal any of the data.

Risks Associated with Internet Access

Although the Internet has brought computer crime, such as cracking, to the public's attention, network security is more often compromised "from the inside" than from external sources. Nevertheless, the threat of outside intruders is very real, and it will only grow as more people gain access to the Internet.

At the same time, users need to be careful when they connect to the Internet. Even the most popular Web browsers sometimes contain bugs in their most recent releases that permit scripts to access your system while you're connected to the Internet, potentially for the purpose of causing damage. And be careful what information you provide while browsing the Web. Some sites will capture that information to use when attempting to break into systems. Bear in mind that crackers are creative and typically revel in devising new ways of breaking into systems. As a result, new Internet-related security threats arise frequently. By keeping your software current, staying abreast of emerging security threats, and designing your Internet access wisely, you can prevent most of these threats. Common Internet-related security breaches include the following:

- A firewall may not be adequate protection, if it is configured improperly. For example, it may allow outsiders to obtain internal IP addresses, then use those addresses to pretend that they have authority to access your internal network from the Internet—a process called **IP spoofing**. Alternately, a firewall may not be configured correctly to perform even its simplest function—preventing unauthorized packets from entering the LAN from outside. (You will learn more about firewalls later in this chapter.) Correctly configuring a firewall is one of the best means to protect your internal LAN from Internet-based attacks.
- When a user Telnets or FTPs to your site over the Internet, his or her user ID and password will be transmitted in plain text—that is, unencrypted. Anyone monitoring the network can pick up the user ID and password and use it to gain access to the system.
- Crackers may obtain information about your user ID from newsgroups, mailing lists, or forms you have filled out on the Web (for example, to register to win a new car on a promotional site).
- While users remain logged onto Internet chat sessions, they may be vulnerable to other Internet users who might send commands to their machines that cause the screen to fill with garbage characters and require them to terminate their chat sessions. This type of attack is called **flashing**.

- After gaining access to your system through the Internet, a cracker may launch denial-of-service attacks. A **denial-of-service attack** occurs when a system becomes unable to function because it has been deluged with data transmissions or otherwise disrupted. This incursion is a relatively simple attack to launch (for example, a cracker could create a looping program that sent thousands of e-mail messages to your system per minute). The easiest resolution of this problem is to bring down the attacked server, then reconfigure the firewall to deny service (in return) to the attacking machine. Denial-of-service attacks may also result from malfunctioning software. In Chapter 13, you learned how to apply patches to your server's operating system and utilities and research vendors' update alerts. Regularly performing these upgrades is essential to maintaining network security.

ADDRESSING RISKS ASSOCIATED WITH PEOPLE

As you have learned, most network security breaches occur from within an organization, and many take advantage of human errors. This section describes how to minimize the risk of break-ins by communicating with and managing the users in your organization via a thoroughly planned security policy. Before any hardware or software measures can offer effective protection, a security policy must be implemented that tells users how to set secure passwords and that makes critical data accessible only to authorized personnel.

An Effective Security Policy

The first step in securing your network is devising and implementing a security policy. This document identifies your security goals, risks, levels of authority, designated security coordinator and team members, the responsibilities for each team member, and the responsibilities for each employee. In addition, it specifies how to address security breaches. It should not state exactly which hardware, software, architecture, or protocols will be used to ensure security, nor how hardware or software will be installed and configured. These details will change from time to time and should be shared only with authorized network administrators or managers.

Security Policy Goals

Before drafting a security policy, you should understand why the security policy is necessary and how it will serve your organization. Typical goals for security policies are as follows:

- Ensuring that authorized users have appropriate access to the resources they need
- Preventing unauthorized users from gaining access to the network, systems, programs, or data

- Protecting sensitive data from unauthorized access, both from within and from outside the organization
- Preventing accidental damage to hardware or software
- Preventing intentional damage to hardware or software
- Creating an environment where the network and systems can withstand and, if necessary, quickly respond to and recover from any type of threat
- Communicating each employee's responsibilities with respect to maintaining data integrity and system security

A company's security policy may also include content that does not pertain to computers or networks. For example, it might state that each employee must shred paper files that contain sensitive data or that each employee is responsible for signing in his or her visitors at the front desk and obtaining a temporary badge for them. Non-computer-related aspects of security policies are beyond the scope of this chapter, however.

After defining the goals of your security policy, you can devise a strategy to attain them. First, you might form a committee composed of managers and interested parties from a variety of departments, in addition to your network administrators. The more decision-making people you can involve, the more supported and effective your policy will be. This committee can assign a security coordinator, who will then drive the creation of a security policy.

To increase the acceptance of your security policy in your organization, tie security measures to business needs and clearly communicate the potential effects of security breaches. For example, if your company sells clothes over the Internet and a two-hour outage (as could be caused by a cracker who uses IP spoofing to gain control of your systems) could cost the company $1 million in lost sales, make certain that users and managers understand this fact. If they do, they will be more likely to embrace the security policy.

A security policy must address an organization's specific risks. To understand your risks, you should conduct a security audit that identifies vulnerabilities and rates both the severity of each threat and its likelihood of occurring, as described earlier in this chapter. Once risks are identified, the security coordinator should assign one person the responsibility for addressing that threat.

For example, imagine that you are the network administrator for a nonprofit organization that collects blood donations from the public and arranges to ship them to victims of disasters across the country. Your LAN contains not only your organization's financial and personnel data, but also databases listing all blood donors in your area and the last date that they gave blood. In addition, your network contains records of the people whom

your organization has assisted during the last five years. Your network is connected through the Internet to other, similar organizations, so that you can share your resources with them and they can help you with your needs. What security risks exist in this system, and how should you manage them?

First, your servers hold a great deal of potentially sensitive information—not only financial information, but also health and community data. To prevent unauthorized access to this information, one person should be assigned the task of protecting it. Second, your network presumably has a number of workstations attached to it. At least one person should be given the responsibility of making sure that PCs do not hold sensitive data on their hard disks and that PC users understand and adhere to security policies. Third, your network has a link to the Internet. One person should therefore be assigned the task of verifying that the Internet connection does not permit crackers to access your internal network.

Security Policy Content

After your risks are identified and responsibilities for managing them are assigned, the policy's outline should be generated with those risks in mind. Some subheadings for the policy might include the following: Password policy; Software installation policy; Confidential and sensitive data policy; Network access policy; E-mail use policy; Internet use policy; Modem use policy; Remote access policy; Policies for connecting to remote locations, the Internet, and customers' and vendors' networks; Policies for use of laptops and loaner machines; and Computer room access policy. Although compiling all of this information might seem like a daunting task, the process will ensure that everyone understands the organization's stance on security and the reasons why it is so important.

The security policy should clearly explain to users what they can and cannot do and how these measures protect the network's security. Clear and regular communication about security policies will make them more acceptable and better understood. One idea for making security policies more sustainable is to distribute a "security newsletter" that keeps security issues fresh in everyone's mind. Perhaps the newsletter could highlight industry statistics about significant security breaches and their effect on the victimized organizations. You might also hold a contest for guessing a password, thereby demonstrating how unsafe passwords threaten security.

Another tactic is to create a separate section of the policy that applies only to users. Within the users' section, divide security rules according to the particular function or part of the network to which they apply. This approach will make the policy easier for users to read and understand; it will also prevent them from having to read through the entire document. The following abridged example shows a section organized in this way:

3.2.1 Passwords
Users may not share passwords with friends or relatives.
Users must choose passwords that exceed six characters and are composed of both letters and numbers.
Users should choose passwords that bear no resemblance to a spouse's name, pet's name, birth date, anniversary, or other widely available information.
Users must change their passwords every 90 days.
Users may not use the same password until one year after its expiration.
Users may not write down their passwords or send them in e-mail correspondence.
3.2.2 Networks
Users must ensure that confidential data transmitted over the network are encrypted.
Users should be aware that most e-mail programs allow administrators to read any messages that are sent from or received by the system.
Users may not have modems attached to their machines unless they have received authorization from their supervisor.
Users may not install software obtained from the Internet or other outside sources.

Notice that this sample policy asks users to encrypt confidential data when transmitting such information over the network. But how will users know what type of information is confidential? Your security policy should define what "confidential" means to your organization. In general, information is confidential if it could be used by other parties to impair your organization's functioning, decrease your customers' confidence, cause a financial loss, damage your organization's status, or give a significant advantage to a competitor. If you work in an environment such as a hospital, where most data are sensitive or confidential, however, your security policy should classify information in degrees of sensitivity that correspond to how strictly its access is regulated. For example, "top-secret" data may be accessible only by the organization's CEO and vice presidents, whereas "confidential" data may be accessible only to those who must modify or create it (for example, doctors or hospital accountants).

Response Policy

Finally, a security policy should provide for a planned response in the event of a security breach. The response policy should identify the members of a response team, all of whom should clearly understand the security policy, risks, and measures in place. Each team member should be assigned a role and responsibilities. Like a disaster recovery response team, the security response team should regularly rehearse their defense by participating in a security threat drill. Some suggestions for team roles are listed below:

- *Dispatcher*— This team member is the person on call who first notices or is alerted to the problem. The dispatcher notifies the lead technical support specialist and then the manager. He or she opens a record on the incident, detailing the time it began, its symptoms, and any other pertinent information

about the situation. The dispatcher remains available to answer calls from clients or employees or to assist the manager.

- *Manager*—This team member coordinates the resources necessary to solve the problem. If in-house technicians cannot handle the break-in, the manager should find outside assistance. The manager also ensures that the security policy is followed and that everyone within the organization is aware of the situation. As the response ensues, the manager continues to monitor events and communicate with the public relations specialist. After the incident has been resolved, the manager should hold a postmortem meeting to discuss how the breach happened, how the problem was resolved, and what measures are being taken to prevent a recurrence.
- *Technical support specialist*—This team member focuses on only one thing: solving the problem as quickly as possible. After the situation has been resolved, the technical support specialist describes in detail what happened and assists the manager in finding ways to avert such an incident in the future. Depending on the size of the organization and the severity of the incident, this role may be filled by more than one person.
- *Public relations specialist*—If necessary, this team member learns about the situation and the response, then acts as official spokesperson for the organization to the public.

After resolving a problem, you need to review what happened, determine how it might have been prevented, then implement those measures to prevent future problems. Better than having to learn from your own mistakes, though, is learning from others' mistakes. By searching the Web, you can find many examples of security breaches that cost businesses millions of dollars in lost revenue either because their systems failed or because they lost valuable trade secrets.

Passwords

Choosing a secure password is one of the easiest and least expensive ways to guard against unauthorized access. Unfortunately, too many people prefer to use an easy-to-remember password. If your password is obvious to you, however, it may also be easy for a cracker to figure out. The following guidelines for selecting passwords should be part of your organization's security policy. It is especially important for network administrators to choose difficult passwords, and also to keep passwords confidential and to change them frequently.

Tips for making and keeping passwords secure include the following:

- Do not use familiar information, such as your birth date, anniversary, pet's name, child's name, spouse's name, own name or nickname, user ID, phone number, address, or any other words or numbers that others might associate with you.

- Do not use any word that might appear in a dictionary. Crackers can use programs that try a combination of your user ID and every word in a dictionary to gain access to the network.
- Make the password longer than six characters—the longer, the better.
- Choose a combination of letters and numbers; add special characters, such as exclamation marks or hyphens, if allowed.
- Do not write down your password or share it with others.
- Change your password at least every 90 days, or more frequently, if desired. If you are a network administrator, establish controls through the network operating system to force users to change their passwords at least every 90 days. If you have access to sensitive data, change your password even more frequently.

Password guidelines should be clearly communicated to everyone in your organization through your security policy. Although users may grumble about having to choose a combination of letters and numbers and change their passwords frequently, you can assure them that the company's financial and personnel data will be safer as a result. No matter how much your colleagues protest, do not back down from your password requirements. Many companies mistakenly require employees only to use a password, without helping them choose a good one. This oversight increases the risk of security breaches.

PHYSICAL SECURITY

Another important element in network security is the restriction of physical access to its components. At the very least, only authorized networking personnel should have access to computer rooms. If computer rooms are not locked, intruders may easily steal equipment or sabotage software and hardware. For example, a malicious visitor could slip into an unsecured computer room and take control of a NetWare server console where an administrator is logged in, then shut down the machine—or worse, reformat its hard disk. Although a security policy may define who has access to the computer room, locking the computer room is necessary to keep unauthorized individuals out.

It isn't only the computer room that must be secured. Think of all the points at which your systems or data could be compromised: hubs or switches in a wiring closet, an unattended workstation at someone's desk, a telecommunications closet where your leased line to the Internet terminates, a storage room for archived data and backup tapes. If a wiring closet is left unlocked, for example, a prankster could enter, grab a handful of wires, and pull them out of the patch panels.

Locks may be either physical or electronic. Many large organizations require authorized employees to wear electronic access badges. These badges can be programmed to allow their owner access to some, but not all, rooms in a building. Figure 15-1 depicts a typical badge access security system.

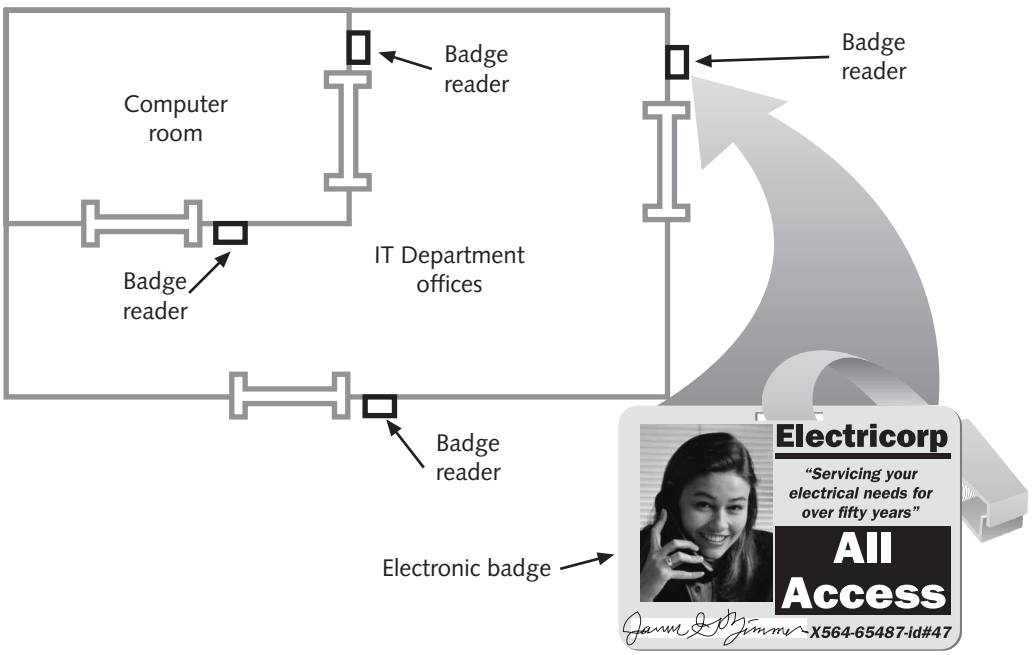

Figure 15-1 A badge access security system

A less-expensive alternative to the electronic badge access system consists of locks that require entrants to punch a numeric code to gain access. For added security, these electronic locks can be combined with key locks. A more expensive solution involves **bio-recognition access**, in which a device scans an individual's unique physical characteristics, such as the color patterns in her eye's iris or the geometry of her hand, to verify her identity. On a larger scale, organizations may regulate entrance through physical barriers to their campuses, such as gates, fences, walls, or landscaping.

Many IT departments also use closed-circuit TV systems to monitor activity in secured rooms. Surveillance cameras may be placed in computer rooms, telco rooms, supply rooms, and data storage areas, as well as facility entrances. A central security office may display several camera views at once, or it may switch from camera to camera. The footage generated from these cameras is usually saved for a time in case it's needed in a security breach investigation or prosecution.

As with other security measures, the most important way to ensure physical security is to plan for it. You can begin your planning by asking questions related to physical security checks in your security audit. Relevant questions include the following:

- Which rooms contain critical systems or data and need to be secured?
- Through what means might intruders gain access to the facility, computer room, telecommunications room, wiring closet, or data storage areas (includ-

ing not only doors, but also windows, adjacent rooms, ceilings, temporary walls, hallways, and so on)?

- How and to what extent are authorized personnel granted entry? (Do they undergo background or reference checks? Is their need for access clearly justified? Are their hours of access restricted? Who ensures that lost keys are reported?)
- Are employees instructed to ensure security after entering or leaving secured areas (for example, by not propping open doors)?
- Are authentication methods (such as ID badges) difficult to forge or circumvent?
- Do supervisors or security personnel make periodic physical security checks?
- Are all combinations, codes, or other access means to computer facilities protected at all times, and are these combinations changed frequently?
- Do you have a plan for documenting and responding to physical security breaches?

ADDRESSING RISKS ASSOCIATED WITH HARDWARE AND DESIGN

Addressing the risks associated with people is just one part of a comprehensive security approach. Even if you restrict access to computer rooms, teach employees how to select secure passwords, and enforce a security policy, breaches may still occur due to poor LAN or WAN design. In this section, you will learn how to address some security risks via intelligent networking hardware and design.

Of course, the optimal way to prevent external security breaches from affecting your LAN is to not connect your LAN to the outside world at all. This option is impractical in today's business environment, however. The next best protection is to restrict access at every point where your LAN connects to the rest of the world. This principle forms the basis of hardware- and design-based security.

Much of the information covered in this section builds upon material discussed in Chapter 7, such as WAN design, VPNs, and remote connectivity. It may be helpful to review Chapter 7 before reading this section.

Firewalls

As you learned early in this chapter, a firewall is a specialized device that selectively filters or blocks traffic between networks. A firewall typically involves a combination of hardware and software and may reside between two interconnected private networks or, more typically, between a private network and a public network (such as the Internet), as shown in Figure 15-2. Many types of firewalls exist, and a detailed discussion of each is beyond the scope of this book. To understand secure network design and to qualify for

Net+ certification, however, you should recognize which functions firewalls can provide, where they can appear on a network, and how to decide what you need in a firewall.

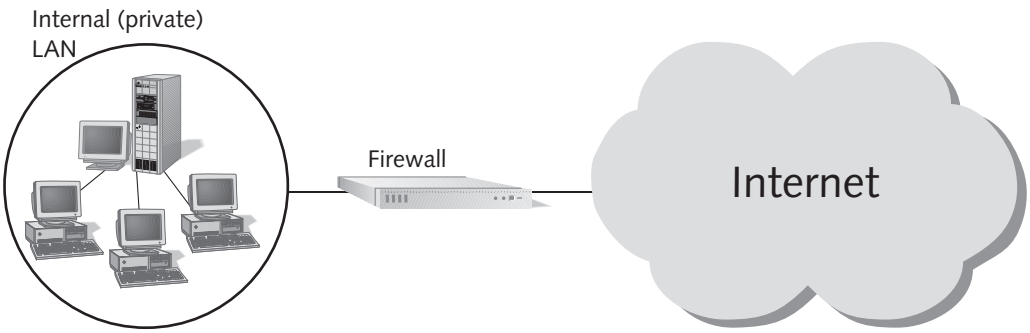

Figure 15-2 Placement of a firewall between a private network and the Internet

The simplest and most common form of a firewall is a **packet-filtering firewall**, which is a router that operates at the Data Link and Transport layers of the OSI Model. It examines the header of every packet of data that it receives to determine whether that type of packet is authorized to continue to its destination. Packet-filtering firewalls are also called **screening firewalls**. An example of a popular packet-filtering firewall is the Cisco PIX 525, pictured in Figure 15-3.

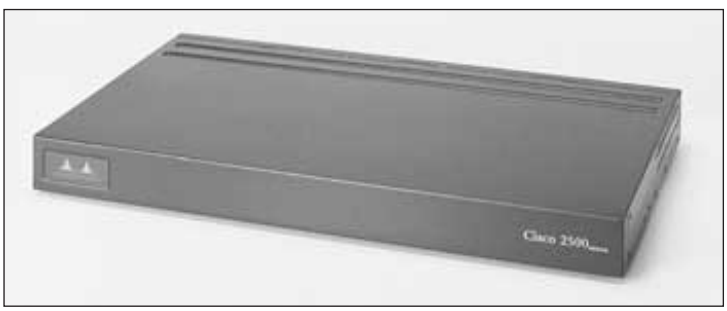

Figure 15-3 A packet-filtering firewall

These types of firewalls require a great deal of custom configuration to be effective; that is, the network administrator must configure the firewall to accept or deny certain types of traffic. Some of the criteria that a firewall might use to accept or deny data include the following:

- Source and destination IP addresses
- Source and destination ports (for example, ports that supply TCP/UDP connections, FTP, Telnet, SNMP, RealAudio, and so on)
- The TCP, UDP, or ICMP protocols

- A packet's status as the first packet in a new data stream or a subsequent packet
- A packet's status as inbound to or outbound from your private network
- A packet's status as originating from or being destined for an application on your private network

Based on these options, a network administrator could configure her firewall, for example, to prevent any IP address that does not begin with "196.57," the first two octets of the addresses on her network, from accessing the network's router and servers. Furthermore, she could disable—or block—certain well-known ports, such as the FTP ports (20 and 21) through the router's configuration. Blocking ports prevents *any* user from connecting to and completing a transmission through those ports. This technique is useful to further guard against unauthorized access to the network. In other words, even if a cracker were able to spoof an IP address that began with "196.57," he could not access the FTP ports (which are notoriously insecure) on the firewall. Ports can be blocked not only on firewalls, but also on routers, servers, or any device that uses ports. For example, if you established a Web server for testing but did not want anyone in your organization to connect to your Web pages through his or her browsers, you could block port 80 on that server.

You will recognize examples of firewall placement in most VPN architectures. For example, you might design a VPN that uses the Internet to connect your Milwaukee office with your Denver office. To ensure that only traffic from Milwaukee can access your Denver LAN, you could install a packet-filtering firewall between the Denver LAN and the Internet that accepts incoming traffic only from IP addresses that match the IP addresses on your Milwaukee LAN. In a way, the firewall acts like a bouncer at a private club who checks everyone's ID and ensures that only club members enter through the door. In the case of the Milwaukee-Denver VPN, the firewall will discard any data packets that arrive at the Denver firewall and do not contain source IP addresses that match those of Milwaukee's LAN.

In another example, suppose your network in Denver hosts a server that stores confidential employee information, such as payroll and health benefits, which only the Denver-based human resources manager should be able to access. In this situation, you could add a filter in the firewall to block all external traffic (from the Internet as well as from the Milwaukee LAN) from reaching the destination address of that server.

Because you must tailor a firewall to your network's needs, you cannot simply purchase one, install it between your private LAN and the Internet, and expect it to offer much security. Instead, you must first consider what type of traffic you want to filter, then configure the firewall accordingly. It may take weeks to achieve the best configuration—not so strict that it prevents authorized users from transmitting and receiving necessary data, and not so lenient that you risk security breaches. Further complicating the matter is that you may need to create exceptions to the rules. For example, suppose that your human resources manager is working out of the Milwaukee office while recruiting new

employees and needs to access the Denver server that stores payroll information. In this instance, the Denver network administrator might create an exception to allow transmissions from the human resources manager's workstation's IP address to reach that server. In the networking profession, creating an exception to the filtering rules is called "punching a hole" in the firewall.

Because packet-filtering routers operate at the Network and Transport layers of the OSI Model and examine only network addresses, they cannot distinguish between a user who is trying to breach the firewall and a user who is authorized to do so. To ensure that an unauthorized user does not simply sit down at the workstation belonging to an authorized user and try to circumvent the firewall, a more sophisticated technique—such as user authentication—is necessary.

One approach to enhancing the security of the Network and Transport layers provided by firewalls is to combine a packet-filtering firewall (hardware device) with a proxy service. A **proxy service** is a software application on a network host that acts as an intermediary between the external and internal networks, screening all incoming and outgoing traffic. The network host that runs the proxy service is known as a **proxy server**, or gateway. Proxy servers manage security at the Application layer of the OSI Model. To understand how they work, think of the secure data on a server as the president of a country and the proxy server as the secretary of state. Rather than having the president risk his or her safety by leaving the country, the representative travels abroad, speaking for the president and bringing information back to the president. In fact, foreign leaders may never actually meet the dignitary. Instead, the representative acts as his or her proxy.

Although a proxy server appears to the outside world as an internal network server, in reality it is merely another filtering device for the internal LAN. Among other things, it prevents the outside world from discovering the addresses of the internal network. For example, suppose your LAN uses a proxy server, and you want to send an e-mail message from your workstation to your mother via the Internet. Your message would first go to the proxy server (depending on the configuration of your network, you may or may not have to log on separately to the proxy server first). The proxy server would repackage the data frames that make up the message so that, rather than your workstation's IP address being the source, the proxy server would insert its own IP address as the source. Next, the proxy server would pass your repackaged data to the packet-filtering firewall. The firewall would verify that the source IP address in your packets is valid (that it came from the proxy server) and then send your message to the Internet. Examples of proxy server software include Novell's BorderManager and Microsoft's Internet Security and Acceleration (ISA) Server 2000, an optional service for Windows 2000 servers. Figure 15-4 depicts how a proxy server might fit into a WAN design.

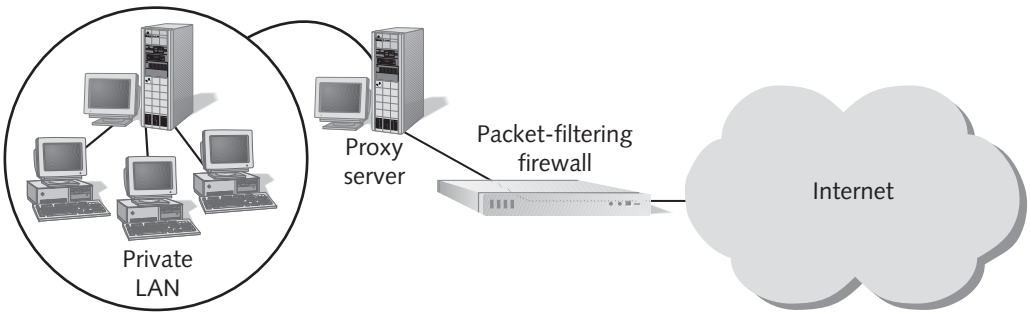

Figure 15-4 A proxy server used on a WAN

Many more sophisticated firewalls—both hardware- and software-based—exist. Choosing the appropriate firewall for your network can be a difficult task. Among the factors you will want to consider when making your decision are the following:

- Does the firewall support encryption? (You will learn more about encryption later in this chapter.)
- Does the firewall support user authentication?
- Does the firewall allow you to manage it centrally and through a standard interface (such as one that uses SNMP)?
- How easily can you establish rules for access to and from the firewall?
- Does the firewall support filtering at the highest layers of the OSI Model, not just at the Data Link and Transport layers?
- Does the firewall provide logging and auditing capabilities, or alert you to possible intrusions?
- Does the firewall protect the identity of your internal LAN's addresses from the outside world?

Remote Access

As you learned in Chapter 7, many companies supply traveling employees, telecommuters, or distant vendors with access to their private LANs or WANs. This type of access is often referred to as remote access. When working with remote access, you must remember that any entry point to a LAN or WAN creates a potential security risk. In other words, if an employee can get to your network in New York from his hotel room in Rome, a smart cracker can likely do the same. You can, however, take advantage of techniques designed to minimize the possibility of such unauthorized remote access. For example, firewalls can prevent certain addresses and users from gaining access to your LAN from the outside. In this section, you will learn about other security measures tailored to remote access solutions, such as remote control and dial-up networking.

Remote Control

Recall from Chapter 7 that remote control systems enable a user to connect to a host system on a network from a distance and use that system's resources as if the user were sitting in front of it. This type of access can have benefits for employees who work at home or who travel frequently. Although such remote control systems can be convenient, they can also present serious security risks. Most remote control software programs (for example, Symantec Corporation's pcAnywhere or Computer Associates International, Inc.'s Control*IT*) offer features that increase the security of remote control systems. If you intend to allow remote control access to a host on your LAN, you should investigate these security features and know how to implement them correctly. Important security features that you should seek in a remote control program include the following:

- A logon ID and password requirement for gaining access to the host system.
- The ability of the host system to call back. This feature enables a remote user to dial into the network, enter a user ID, and hang up. The host system then calls the user back at a predetermined number (the authorized user's modem number), thus preventing a cracker from taking over a system even if he or she obtains the correct user ID and password for the host system.
- Support for data encryption on transmissions between the remote user and the system.
- The ability to leave the host system's screen blank while a remote user works on it. This feature prevents people walking by from seeing (potentially confidential) data that the remote user is accessing.
- The ability to disable the host system's keyboard and mouse. Essentially, this feature turns the host system into a terminal that responds to only remote users.
- The ability to restart the host system when a remote user disconnects from the system. This feature prevents anyone from reviewing what happened during the remote user's session or gaining access if the session was accidentally terminated before the remote user could properly log off.

Dial-up Networking

In Chapter 7, you learned about different ways for remote users to log onto a network. One method involved having users dial into a remote access server attached to the network, also known as dial-up networking. Like other remote access solutions, this approach presents security risks. In this section, you will learn how to make dial-up networking more secure.

Dial-up networking differs from remote control in that it effectively turns a remote workstation into a node on the network, through a remote access server. When choosing a remote access software package, you should evaluate its security. A secure remote access server package will include at least the following features:

- Logon ID and password authentication
- The ability to log all dial-up connections, their sources, and their connection times
- The ability to perform callbacks to users who initiate connections
- Centralized management of dial-up users and their rights on the network

In environments where more than a few dozen simultaneous dial-up connections must be supported and their user IDs and passwords managed, a special kind of server known as a **Remote Authentication Dial-In User Service (RADIUS)** may be implemented to offer authentication services to the network's access server (which may run Windows 2000's RAS or Novell's NAS, for example). RADIUS provides a single, centralized point of authentication for dial-in users. It is highly scalable, as it can attach to pools containing hundreds of modems. In addition, RADIUS is also more secure than using a simple remote access solution because its method of authentication prevents users' IDs and passwords from traveling across the phone line in clear text format.

Many Internet service providers use a RADIUS server to allow their subscribers access to the Internet through their modem pools. Other organizations employ it as a central authentication point for mobile or remote users. Figure 15-5 illustrates these two methods for allowing remote users to connect using RADIUS authentication.

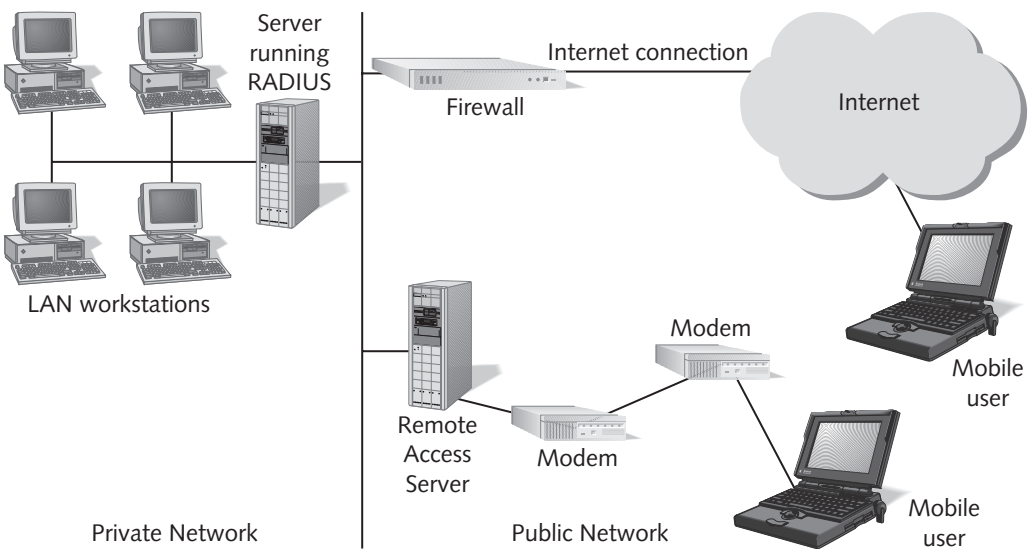

Figure 15-5 A RADIUS server providing central authentication

RADIUS may run on UNIX, Windows 2000, or NetWare networks. A similar, but earlier version of a centralized authentication system is **Terminal Access Controller Access Control System (TACACS)**.

As you learned in Chapter 7, dial-up networking depends on special protocols—for example, PPP or SLIP. Later in this chapter, you will learn about additional protocols that make dial-up networking part of a secure virtual private network.

ADDRESSING RISKS ASSOCIATED WITH PROTOCOLS AND SOFTWARE

Now that you have learned how to design a more secure network, it's time to examine the software tools that can help you protect your network against unauthorized access. Be sure to remember that all of these security techniques complement one another. No single technique is more important than any other, and if you neglect a technique, you may put your network at risk.

The foundation for protecting your organization's data, programs, and access is to master and implement the tools that come with your network operating system. In this section, you will learn the basics of restricting access through this software. You will also learn about a more advanced software-level security technique—encryption.

Network Operating System Authentication

Regardless of whether you run your network on a Novell, Microsoft, or UNIX network operating system, you can implement basic security by restricting what users are authorized to do on a network. This section reiterates what you learned in Chapters 8, 9, and 10 about establishing rights to files and directories on the server.

Every network administrator should understand which resources on the server all users need to access. The rights conferred to all users are called public rights, because anyone can have them and exercising them presents no security threat to the network. In most cases, public rights are very limited. They may include privileges to view and execute programs from the server and to read, create, modify, delete, and execute files in a shared data directory.

In addition, network administrators need to group users according to their security levels and assign additional rights that meet the needs of those groups. As you know, creating groups simplifies the process of granting rights to users. For example, if you work in the IT department at a large college, you will most likely need more than one person to create new user IDs and passwords for students and faculty. Naturally, the staff in charge of creating new user IDs and passwords need the rights to perform this task. You could assign the appropriate rights to each staff member individually, but a more efficient approach is to put all of the personnel in a group, and then assign the appropriate rights to the group as a whole.

In addition to restricting users' access to files and directories on the server, a network administrator can constrain the ways in which users can access the server and its resources. The following is a list of additional restrictions that network administrators can use to strengthen the security of their networks:

- *Time of day*—Some user IDs may be valid only during specific hours—for example, between 8:00 A.M. and 5:00 P.M. Specifying valid hours for an ID can increase security by preventing any ID from being used by unauthorized personnel after hours.
- *Total time logged on*—Some user IDs may be restricted to a specific number of hours per day of logged-on time. Restricting total hours in this way can increase security in the case of temporary IDs. For example, suppose that your organization offers a WordPerfect training class to a group of high-school students one afternoon, and the WordPerfect program and training files reside on your staff server. You might create IDs that could log on for only four hours on that day.
- *Source address*—You can specify that user IDs can log on only from certain workstations or certain areas of the network (that is, domains or segments). This restriction can prevent unauthorized use of logon IDs from workstations outside the network.
- *Unsuccessful logon attempts*—Crackers may repeatedly attempt to log on under a valid ID for which they do not know the password. As the network administrator, you can set a limit on how many subsequent unsuccessful logon attempts from a single user ID the server will accept before blocking that ID from even attempting to log on.

Encryption

Encryption is the use of an algorithm to scramble data into a format that can be read only by reversing the algorithm—that is, by decrypting the data. The purpose of encryption is to keep information private. Many forms of encryption exist, with some being more secure than others. Even as new forms of encryption are developed, new ways of cracking their codes emerge, too.

Encryption is the last means of defense against data theft. In other words, if an intruder has bypassed all other methods of access, including physical security (for instance, he has broken into the telecommunications room) and hardware security (for instance, he has logged onto the router), data may still be safe if they are encrypted. Encryption can protect data when they are stored on a medium, such as a hard disk, or while they are in transit over a communications channel. In order to protect data, encryption provides the following assurances:

- Data were not modified after the sender transmitted them and before the receiver picked them up.

- Data can only be viewed by their intended recipient (or at their intended destination).
- All of the data received at the intended destination were truly issued by the stated sender and not forged by an intruder.

The most popular kind of encryption algorithm weaves a **key** (a random string of characters) into the original data's bits—sometimes several times in different sequences—to generate a unique data block. The scrambled data block is known as **cipher text**. The longer the key, the less easily the cipher text can be decrypted by an unauthorized system. For example, a 512-bit key is considered secure, whereas cipher text generated with a 16-bit key could be cracked in no time.

The process of key encryption is similar to what happens when you finish a card game and place your five-card hand into the deck, then shuffle the deck numerous times. After shuffling, it might take you a while to retrieve your hand. As you can imagine, if you shuffled your five cards into four decks of cards at once, it would be even more difficult to find your original hand. In encryption, theoretically only the computer that is authorized to retrieve the data knows how to unshuffle it and compile the original sequence of data. Figure 15-6 provides a considerably simplified view of key encryption and decryption. Note that actual key encryption does not simply weave a key into the data once, but rather inserts the key, shuffles the data, shuffles the key, inserts another copy of the shuffled key into the shuffled data, shuffles the data again, and so on for several iterations.

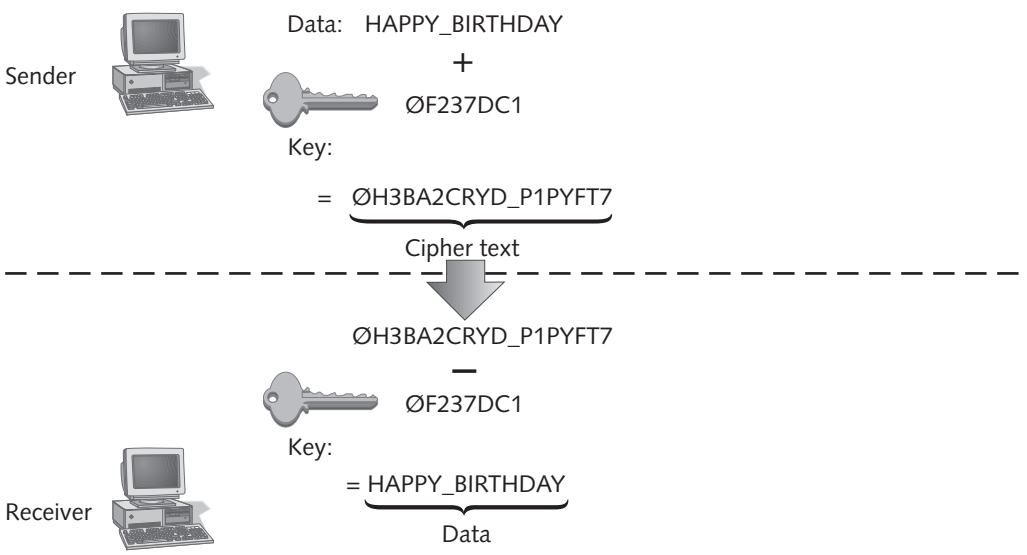

Figure 15-6 Key encryption and decryption

Keys are randomly generated, as needed, by the software that manages the encryption. For example, an e-mail program or a Web browser program may be capable of generating its

own keys to encrypt data. In other cases, special encryption software is used to generate keys. This encryption software works with other types of software, such as word-processing or spreadsheet programs, to encrypt data files before they are saved or transmitted.

Key encryption can be separated into two categories: private key and public key encryption. In **private key encryption** data are encrypted using a single key that only the sender and the receiver know, as depicted in Figure 15-7. This method of key encryption is also known as **symmetric encryption**, because the same key is used during both the transmission and reception of the data. The most popular private key encryption is the **data encryption standard (DES)**, which was developed by IBM in the 1970s.

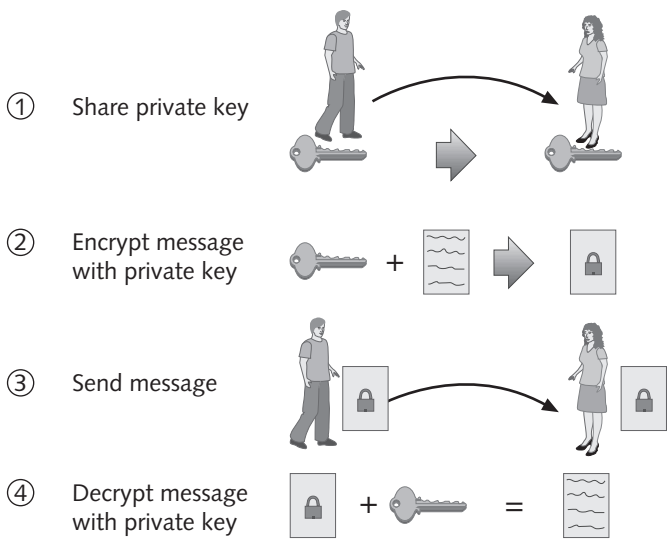

Figure 15-7 Private key encryption

In private key encryption, for example, before Mary can decrypt a message that John sends, he must share his private key with her. Once Mary receives John's encrypted message, she uses a decryption program plus John's private key to decipher the message. The problem with private key encryption is that the sender must somehow share his key with the recipient. For example, John could call Mary and tell her his key, or he could send it to her in an e-mail message. But neither of these methods is very secure. In order to overcome this potential vulnerability, a method of associating publicly available keys with private keys was developed. This method is called public key encryption.

In **public key encryption**, data are encrypted using two keys: one is a key known only to a user (that is, a private key), and the other is a public key associated with the user. A user's public key can be obtained the old-fashioned way—by asking that user—or it can be obtained from a third-party source, such as a public-key server. A **public-key server** is a publicly accessible host (such as an Internet host) that freely provides a list of users'

public keys, much as a telephone book provides a list of peoples' phone numbers. When a user receives a message encrypted with his public key, the recipient's software (for example, his e-mail program) prompts to enter his private key in order to decrypt the message. In other words, the public key has an association with the private key, and a message that has been encrypted with a user's public key can only be decrypted with his private key. The combination of the public key and private key is known as a **key pair**. In this arrangement, every user has a key pair, but one key is known only to the user while the other key is known to those with whom she exchanges data. Because the two users use a different combination of keys, public key encryption is also known as **asymmetric encryption**. Figure 15-8 illustrates the process of public key encryption.

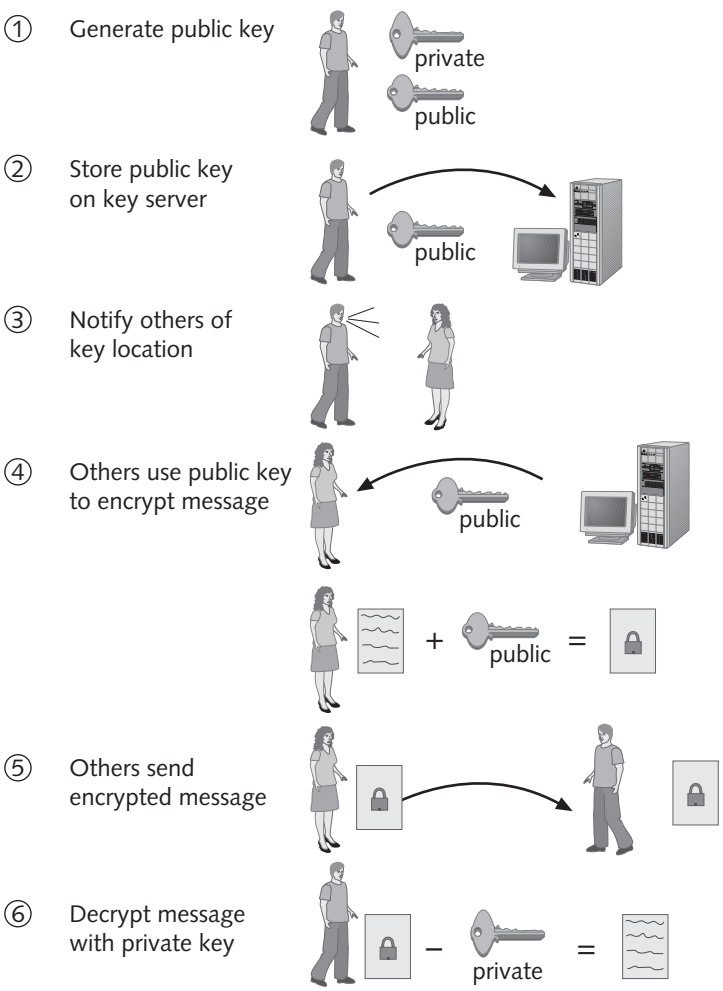

Figure 15-8 Public key encryption

For example, suppose that Mary and John wish to use public key encryption to exchange messages over the Internet. Before sending a message to John, Mary would look up John's public key on a public key server. She would then use her encryption software to scramble her message with John's public key. When John receives the message, his software would recognize that the message has been encoded. Furthermore, the software would recognize that the encryption used John's public key. Based on the public key, it would then prompt John for his private key in order to decrypt the message. Some examples of public key algorithms include RSA (named after its creators, Rivest, Shamir, and Adleman), Diffie-Hellman, and Elliptic-curve cryptography.

With the abundance of private and public keys, not to mention the number of places where each may be kept, users have found a need for easier key management. One answer to this problem is using digital certificates. A **digital certificate** is a password-protected and encrypted file that holds an individual's identification information, including a public key. In the context of digital certificates, the individual's public key is used to verify the sender's digital signature. For example, on the Internet, certificate authorities such as VeriSign, will, for a fee, keep your digital certificate on their server and ensure to all who want to send encrypted messages to you (for example, an order via your e-commerce site) that the certificate is indeed yours. Digital certificates are used in some of the encryption methods discussed below, such as PGP and SSL.

The following sections detail specific public and private key methods of encrypting data as they are transmitted over a network.

Kerberos

Kerberos is a cross-platform authentication protocol that uses key encryption to verify the identity of clients and to securely exchange information once a client logs onto a system. It is an example of a private key encryption service. Kerberos provides significant security advantages over simple network operating system authentication. Whereas an NOS client/server logon process assumes that clients are who they say they are and only verifies a user's name against the password in the NOS database, Kerberos does not automatically trust clients. Instead, it requires a client to prove its identity through a third party. This is similar to what happens when you apply for a passport. The government does not simply believe that you are "Mary Smith," but instead requires you to present proof, such as your birth certificate. In addition to checking the validity of a client, Kerberos communications are encrypted and unlikely to be deciphered by any device on the network other than the client. Contrast this type of transmission to the normally unencrypted and vulnerable communication between an NOS and a client.

In order to understand specifically how a client uses Kerberos, you need to understand some of the terms used when discussing this protocol. In Kerberos terminology, the server that issues keys to clients during initial client authentication is known as the **key distribution center (KDC)**. In order to authenticate a client, the KDC runs an **authentication service (AS)**. An AS issues a **ticket**, which is a temporary set of credentials that a client uses to prove that its identity has been validated (note that a ticket

is not the same as a key, which is used to initially validate its identity). A Kerberos client, or user, is known as a **principal**.

Now that you have learned the terms used by Kerberos, you can follow the process it requires for client/server communication. Bear in mind that the purpose of Kerberos is to connect a valid user with the service that that user wishes to access. In order to accomplish this, both the user and the service must have keys registered with the authentication service. When a user, or principal, wants to access that service, he first logs onto the KDC over the network (on a Windows 2000 network, the KDC is the user's domain controller). Suppose the principal is John Smith, and the service is called "inventory." After logging on, John Smith attempts to log onto the inventory service, thereby, in effect, issuing a message to the authentication service on the KDC that says, "User John Smith wishes to access "inventory." The KDC confirms that John Smith is in its database. Then the AS running on the KDC randomly generates two copies of a new key, called the **session key**. The AS then issues one copy to John Smith and the other copy to the inventory service. Further, it creates a ticket that will allow John Smith to use the inventory service. This ticket contains the inventory service key and can only be decrypted by John Smith's key. The AS sends the ticket to John Smith. John Smith's computer decrypts the session key with John Smith's personal key. It then creates a timestamp associated with his request, and encrypts this timestamp with the session key. The encrypted timestamp is known as the **authenticator**. This timestamp will help the service verify that the ticket is indeed associated with John Smith's request to use the inventory service. Next, John Smith's computer sends his ticket and authenticator to the service. The service decrypts the ticket using its own key and decrypts the authenticator using its session key. Finally, the service has verified that the principal requesting its use is truly John Smith, as the KDC indicated.

The events described above illustrate the original version of the Kerberos authentication process. The problem with the original version was that a user would have to request a separate ticket each time he wished to use a different service. To alleviate this inconvenience, Kerberos developers created the **ticket granting service (TGS)**, an application separate from the AS that also runs on the KDC. So that the client does not need to request a new ticket from the TGS each time it wants to use a different service on the network, the TGS issues the client a **ticket granting ticket (TGT)**. After receiving the TGT, any time that the user wishes to contact a service, he requests a ticket not from the AS, but from the TGS. Furthermore, the reply is encrypted not with the user's personal key, but with the session key that the AS provided for use with the TGS. Inside that reply is the new session key for use with the regular service. The rest of the exchange continues as described above.

Kerberos, which is named after the three-headed dog in Greek mythology who guarded the gates of Hades, was designed at Massachusetts Institute of Technology (MIT). MIT still provides free copies of the Kerberos code. In addition, many software vendors have developed their own versions of Kerberos.

Pretty Good Privacy (PGP)

You have probably exchanged e-mail messages over the Internet without much concern for what happens with your message between the time you send it and when your intended recipient picks it up. In addition, you have probably picked up e-mails from friends without thinking that they might not be from your friends, but rather from other users who are impersonating your friends over the Internet. In fact, typical e-mail communication is a highly insecure form of data exchange. The contents of a message are usually sent in clear (that is, unencrypted) text, which makes it readable by anyone who can capture the message on its way from you to your recipient. In addition, a person with malicious intentions can easily pretend they are someone else. For example, if your e-mail address is joe@trinketmakers.com, someone else could assume your address and send messages that appear to be sent by joe@trinketmakers.com. In order to secure e-mail transmissions, a computer scientist named Phil Zimmerman developed PGP in the early 1990s. **Pretty Good Privacy (PGP)** is a public key encryption system that can verify the authenticity of an e-mail sender and encrypt e-mail data in transmission. PGP is freely available as both an open source and a proprietary software package. Since its release, it has become the most popular tool for encrypting e-mail.

Secure Sockets Layer (SSL)

Secure Sockets Layer (SSL) is a method of encrypting TCP/IP transmissions—including Web pages and data entered into Web forms—en route between the client and server using public key encryption technology. If you trade stocks or purchase goods on the Web, for example, you are most likely using SSL to transmit your order information. SSL is popular in part because it is widely accepted. The most recent versions of Web browsers such as Netscape Communicator and Internet Explorer include SSL client support in their software.

If you have used the Web, you have probably noticed that URLs for most Web pages begin with the HTTP prefix, which indicates that the request will be handled by TCP/IP port 80 using the HTTP protocol. When Web page URLs begin with the prefix **HTTPS** they are requiring that their data be transferred from server to client and vice versa using SSL encryption. HTTPS uses the TCP port number 443, rather than port 80. Once an SSL connection has been established between a Web server and client, the client's browser indicates this by showing a padlock in the lower-right corner of the screen (this applies to Internet Explorer and Netscape Communicator versions 4.0 and higher).

Each time a client and server establish an SSL connection, they also establish a unique **SSL session**, or an association between the client and server that is defined by an agreement on a specific set of encryption techniques. An SSL session allows the client and server to continue to exchange data securely as long as the client is still connected to the server. An SSL session is created by the SSL handshake protocol, one of several protocols within SSL, and perhaps the most significant. As its name implies, the **handshake protocol** allows the client and server to authenticate (or introduce) each other and establishes terms for how they will securely exchange data. For example, when you are

connected to the Web and you decide to open your bank's account access URL, your browser initiates an SSL connection with the handshake protocol. The handshake protocol sends a special message to the server, called a **client_hello** message, which contains information about what level of security your browser is capable of accepting and what type of encryption your browser can decipher (for example, RSA or Diffie-Hellman). The client_hello message also establishes a randomly generated number that uniquely identifies your client and another number that identifies your SSL session. The server responds with a **server_hello** message that confirms the information it received from your client and agrees to certain terms of encryption based on the options your client supplied. Depending on the Web server's preferred encryption method, the server may choose to issue your browser a public key or a digital certificate at this time. Once the client and server have agreed on the terms of encryption, they will begin exchanging data.

SSL was originally developed by Netscape. Since that time, the Internet Engineering Task Force (IETF) has attempted to standardize SSL in a protocol called **Transport Layer Security (TLS)**. Besides standardizing SSL for use with software from multiple vendors, IETF also aims to create a version of SSL that will encrypt UDP as well as TCP transmissions. TLS, which will likely be supported by new Web browsers, uses slightly different encryption algorithms than SSL, but otherwise is very similar to the most recent version of SSL.

Internet Protocol Security (IPSec)

The **Internet Protocol Security (IPSec)** protocol defines encryption, authentication, and key management for TCP/IP transmissions. It is an enhancement to IPv4 and is native to the newer, IPv6 standard. IPSec is somewhat different from other methods of securing data in transit. Rather than applying encryption to a stream of data, IPSec actually encrypts data by adding security information to the header of all IP packets. In effect, IPSec transforms the data packets. To do so, IPSec operates at the Network layer (Layer 3) of the OSI Model.

IPSec accomplishes authentication in two phases. The first phase is key management and the second phase is encryption. **Key management** refers to the way in which two nodes agree on common parameters for the keys they will use. IPSec relies on **Internet Key Exchange (IKE)** for its key management. IKE is a service that runs on UDP port 500. Once IKE has established the rules for the type of keys two nodes will use, IPSec invokes its second phase, encryption. In this phase, two types of encryption may be used: **authentication header (AH)** and **encapsulation security payload (ESP)**. It is not important to know the inner workings of these services in order to qualify for Network+ certification, but you should be aware that both types of encryption provide authentication of the IP packet's data payload through public key techniques. In addition, EPS encrypts the entire IP packet for added security.

IPSec can be used with any type of TCP/IP transmission. However, it most commonly runs on routers or other connectivity devices in the context of VPNs. Because VPNs are

used to transmit private data over public networks, they require strict encryption and authentication to ensure that data are not compromised. The next section focuses on traditional VPN security measures.

VIRTUAL PRIVATE NETWORK (VPN) SECURITY

As you learned in Chapter 7, virtual private networks (VPNs) are private networks that use public channels to connect clients and servers. Often VPNs integrate a wide variety of clients, from dial-up users at home to networked workstations in offices to Web servers at an ISP. The mix of client types, transmission methods, and services used by VPNs adds to their design complexity, as well as to the complexity of their security needs. Security considerations must be woven into both hardware/design and software for VPNs. These types of networks are so varied and potentially complicated that fully describing their nuances is beyond the scope of this book. In this section, however, you will learn about the significant security techniques particular to VPNs.

VPNs typically use the Internet to connect multiple sites; because the Internet is the largest public network in the world, its use presents obvious security hazards. VPNs often take advantage of firewalls and special protocols that encrypt the data transmitted over public connections. The following sections describe some of the special protocols used in VPN connectivity.

As described in Chapter 7, PPP is a dial-in protocol that belongs in the Data Link layer (Layer 2) of the OSI Model and provides datagram transport services over serial and digital communications lines for the TCP/IP, NetBEUI, and IPX/SPX protocols. PPP originated for use with direct dial-in connections to Windows NT RAS servers. The **Point-to-Point Tunneling Protocol (PPTP)** expands on PPP by encapsulating it so that any type of PPP data can traverse the Internet masked as a pure IP transmission. PPTP supports the encryption, authentication, and LAN access services provided by RAS. Instead of users having to dial directly into an access server, however, they can dial into their ISP using PPTP and thereby gain access to their corporate LAN over the Internet.

The process of encapsulating one protocol to make it appear as another type of protocol is known as **tunneling**. Essentially, tunneling makes a protocol fit a type of network that it wouldn't normally match. PPTP is easy to install, is available at no extra cost with Microsoft networking services, and supports multiple kinds of protocols. For these reasons, it is the most popular VPN tunneling protocol in use today.

PPTP is available with both the server and workstation versions of Windows NT and Windows 2000 as part of RAS. You can purchase an upgrade from Microsoft to enable PPTP to work with the Windows 95 Dial-up Networking client. PPTP support is included automatically in the Windows 98 operating system.

Layer 2 Forwarding (L2F) is similar to PPTP in that it is a Layer 2 protocol that provides tunneling for other protocols and can work with the authentication methods used by PPP. The difference between PPTP and L2F lies in the type of encryption that each supports, and the fact that PPTP was developed by Microsoft, and L2F was developed by Cisco Systems. One disadvantage of L2F as compared to PPTP is that the former protocol requires special hardware on the host system end, whereas PPTP will work with any Windows NT or 2000 server. On the other hand, L2F can encapsulate protocols to fit more than just the IP format, unlike PPTP.

Both PPTP and L2F, however, will gradually be replaced by a third type of tunneling protocol called **Layer 2 Tunneling Protocol (L2TP)**. This Layer 2 tunneling protocol was developed by a number of industry consortia. L2TP is an enhanced version of L2F that, like L2F, supports multiple protocols. Unlike L2F, however, L2TP does not require costly hardware upgrades to implement. It is also optimized to work with the next generation of IP (IPv6) and IPSec.

CHAPTER SUMMARY

- A hacker is someone who masters the inner workings of operating systems and utilities in an effort to better understand them. A cracker is someone who uses his or her knowledge of operating systems and utilities to intentionally damage or destroy data or systems.
- The root is a highly privileged user ID that has all rights to create, delete, modify, move, read, write, or execute files on a system. This term may specifically refer to the administrator on a UNIX-based network. Getting the root ID and password on one system often allows crackers to gain access to attached systems.
- Authentication is the process of verifying a user's validity and authority to use a system. You are familiar with the user ID and password combination. Systems may also base authentication on digital signatures, IP addresses, session IDs, or a combination of these methods. Generally, the more information required for authentication, the more secure the system.
- Every organization should assess its security risks by conducting a security audit, at least annually and preferably quarterly. For each threat, your security audit should rate the severity of its potential consequences, as well as its likelihood.
- One of the most common methods by which an intruder gains access to a network is to simply ask a user for his or her password. This strategy is commonly called social engineering, because it involves manipulating social relationships to gain access.
- Security risks associated with people include the following: intruders or attackers using social engineering to obtain user passwords; an administrator incorrectly creating or configuring user IDs, groups, and their associated rights on a file server; network administrators overlooking security flaws in topology or hardware configu-

ration; network administrators overlooking security flaws in operating system or application configuration; lack of proper documentation and communication of security policies; dishonest or disgruntled employees abusing their file and access rights; a computer or terminal being left logged onto the network while its operator is away; users, or even administrators, choosing easy-to-guess passwords; authorized staff leaving computer room doors propped open or unlocked, thereby allowing unauthorized individuals to enter; and administrators neglecting to remove access and file rights for employees who have left the organization.

- Risks inherent in network hardware and design include the following: twisted-pair cabling that emits electromagnetic radiation; wireless and wire-based transmissions, which can typically be intercepted (transmissions over fiber-based networks cannot); networks that use leased public lines, which are subject to eavesdropping; network hubs that broadcast traffic over the entire segment, thus making transmissions more widely vulnerable to sniffing; unused hub, router, or server ports that can be exploited and accessed by crackers if not disabled; a router's configuration port, accessible by Telnet, that may not be adequately secured; routers that may not be properly configured to mask internal subnets; modems attached to network devices that may be configured to accept incoming calls; dial-in access servers used by telecommuting or remote staff that may not be carefully secured and monitored; and computers hosting very sensitive data that may coexist on the same subnet with computers open to the general public.
- Some risks pertaining to networking protocols and software include the following: TCP/IP security flaws; trust relationships between one server and another; network operating system software "backdoors" or security flaws; a network operating system that allows server operators to exit to a command prompt; administrators who accept default operating system security; and transactions that take place between applications left open to interception.
- A denial-of-service attack occurs when a system becomes dysfunctional because it is deluged with traffic. It is a relatively simple attack to launch, and the easiest resolution is to turn off the attacked server.
- The first step in securing your network should be to devise and implement an enterprise-wide security policy. This document identifies your security goals, risks, levels of authority, designated security coordinator and team members, responsibilities for each team member, responsibilities for each employee, and strategies for addressing security breaches. It should not include specific information on what hardware, software, architecture, or protocols will be used to ensure security, nor should it indicate how hardware or software will be installed and configured.
- Goals for an effective security policy should include the following: ensuring that authorized users have appropriate access to the resources they need, preventing unauthorized users from gaining access to the network and its resources, protecting sensitive data from unauthorized access, preventing accidental damage to hardware or software, preventing intentional damage to hardware or software, creating an environment where the network and systems can withstand and quickly recover

from any type of threat, and communicating each employee's responsibilities with respect to maintaining data integrity.

- Choosing secure passwords is one of the easiest and least expensive ways to guard against unauthorized access. The following guidelines for selecting passwords should be part of your organization's security policy: do not use the familiar types of passwords; do not use any word that can be found in a dictionary; make the password longer than six characters; choose a combination of letters and numbers; add special characters, such as exclamation marks or hyphens, if allowed; do not write down your password or share it with others; change your password at least every 90 days; and, if you are a network administrator, establish controls through the network operating system to force users to change their passwords at least every 90 days.
- One way to help keep a network secure is to restrict access to its physical components. At the very least, computer rooms should allow access only to authorized networking personnel. If computer rooms or wiring closets remain unlocked, intruders may easily enter and steal equipment, or sabotage software and hardware.
- A firewall is a specialized device (typically a router, but possibly only a PC running special software) that selectively filters or blocks traffic between networks. It may be placed between two interconnected private networks or, more typically, between a private network and a public network (such as the Internet).
- The simplest and most common form of firewall is a packet-filtering firewall. This router operates at the Data Link and Transport layers of the OSI Model, examining the header of every packet of data that it receives to determine whether that type of packet is authorized to continue to its destination. Packet-filtering firewalls are also called screening firewalls.
- A more sophisticated security technique is necessary to perform user authentication. One approach is to combine a packet-filtering firewall with a proxy service—a software application on a network host that acts as an intermediary between the external and internal networks, screening all incoming and outgoing traffic.
- The network host that runs the proxy service is known as a proxy server or gateway. Although a proxy server appears to the outside world as an internal network server, in reality it is merely another filtering device for the internal LAN. Among other things, it prevents the outside world from discovering the addresses of the internal network.
- Important security features that you should seek in a remote control program include the following: a login ID and password requirement to gain access to the host system, the ability for the host system to call back, support for data encryption on transmissions between the remote user and the system, the ability to leave the host system's screen blank while a remote user works on it, the ability to disable the host system's keyboard and mouse, and the ability to restart the host system when a remote user disconnects from the system.

- A secure remote access server package will include at least the following features: login ID and password authentication; the ability to log all dial-up connections, their sources, and their connection times; the ability to perform callbacks to users who initiate connections; and centralized management of dial-up users and their rights on the network.
- In environments where more than a few dozen simultaneous dial-up connections must be supported and their user IDs and passwords managed, a special kind of server, known as a Remote Authentication Dial-In User Service (RADIUS), may be implemented to offer authentication services to the network's access server. RADIUS provides a single, centralized point of authentication for dial-in users. It is highly scalable because it can attach to pools containing hundreds of modems.
- Every network operating system provides at least some security by allowing you to limit users' access to files and directories on the network. In addition, network administrators can constrain how those with different types of user IDs can use the network by setting restrictions on, for example, time of day, total time logged on, source address, and number of unsuccessful logon attempts.
- Encryption is the use of an algorithm to scramble data into a format that can be read only by reversing the algorithm—or decrypting the data—to keep the information private. Many forms of encryption exist, with some being more secure than others.
- The most popular kind of encryption algorithm weaves a key (a random string of characters) into the original data's bits, sometimes several times in different sequences, to generate a unique data block. The longer the key, the less easily the encrypted data can be decrypted by an unauthorized system.
- Key encryption comes in two forms: public and private key encryption. You should be familiar with at least the following types of encryption: Kerberos, Pretty Good Privacy (PGP), Secure Sockets Layer (SSL), and Internet Protocol Security (IPSec).
- The Point-to-Point Tunneling Protocol (PPTP) expands on PPP by encapsulating it so that any type of PPP data can traverse the Internet masked as a pure IP transmission. PPTP supports the encryption, authentication, and LAN access services provided by RAS. Instead of users having to dial directly into an access server, they can dial into their ISP using PPTP and gain access to their corporate LAN over the Internet.
- PPTP and L2F differ in the type of encryption supported by each and the fact that PPTP was developed by Microsoft and L2F was developed by Cisco Systems. One disadvantage to L2F as compared to PPTP is that the former protocol requires special hardware on the host system end, whereas PPTP will work with any Windows NT or 2000 server. On the other hand, L2F can encapsulate protocols to fit more than just the IP format, unlike PPTP.
- Layer 2 Tunneling Protocol (L2TP) is an enhanced version of L2F that supports multiple protocols, like L2F; unlike L2F, however, L2TP does not require costly hardware upgrades to implement. L2TP is also optimized to work with the next generation of IP (IPv6) and IPSec.

KEY TERMS

asymmetric encryption — A type of encryption (such as public key encryption) that uses a different key for encoding data than is used for decoding the cipher text.

authentication header (AH) — In the context of IPSec, a type of encryption that provides authentication of the IP packet's data payload through public key techniques.

authentication service (AS) — In Kerberos terminology, the process that runs on a key distribution center (KDC) to initially validate a client who's logging on. The authentication service issues session keys to the client and the service the client wants to access.

authenticator — In Kerberos authentication, the user's timestamp encrypted with the session key. The authenticator is used to help the service verify that a user's ticket is valid.

bio-recognition access — A method of authentication in which a device scans an individual's unique physical characteristics (such as the color patterns in his or her eye's iris or the geometry of his or her hand) to verify the user's identity.

cipher text — The unique data block that results when an original piece of data (such as text) is encrypted (for example, by using a key).

client_hello — In the context of SSL encryption, a message issued from the client to the server that contains information about what level of security the client's browser is capable of accepting and what type of encryption the client's browser can decipher (for example, RSA or Diffie-Hellman). The client_hello message also establishes a randomly generated number that uniquely identifies the client plus another number that identifies the SSL session.

cracker — A person who uses his or her knowledge of operating systems and utilities to intentionally damage or destroy data or systems.

data encryption standard (DES) — A popular private key encryption technique that was developed by IBM in the 1970s.

denial-of-service attack — A security attack caused by a deluge of traffic that disables the victimized system.

digital certificate — A password-protected and encrypted file that holds an individual's identification information, including a public key and a private key. The individual's public key is used to verify the sender's digital signature, and the private key allows the individual to log onto a third-party authority who administers digital certificates.

encapsulation security payload (ESP) — In the context of IPSec, a type of encryption that provides authentication of the IP packet's data payload through public key techniques. In addition, ESP also encrypts the entire IP packet for added security.

encryption — The use of an algorithm to scramble data into a format that can be read only by reversing the algorithm—decrypting the data—to keep the information private. The most popular kind of encryption algorithm weaves a key into the original data's bits, sometimes several times in different sequences, to generate a unique data block.

firewall — A specialized device (typically a router, but possibly only a PC running special software) that selectively filters or blocks traffic between networks.

flashing — A security attack in which an Internet user sends commands to another Internet user's machine that cause the screen to fill with garbage characters. A flashing attack will cause the user to terminate his or her session.

hacker — A person who masters the inner workings of operating systems and utilities in an effort to better understand them. A hacker is distinguished from a cracker in that a cracker will attempt to exploit a network's vulnerabilities for malicious purposes.

handshake protocol — One of several protocols within SSL, and perhaps the most significant. As its name implies, the handshake protocol allows the client and server to authenticate (or introduce) each other and establishes terms for how they will securely exchange data during an SSL session.

HTTPS — The URL prefix that indicates that a Web page requires its data to be exchanged between client and server using SSL encryption. HTTPS uses the TCP port number 443, rather than port 80 (the port that normal HTTP uses).

Internet Key Exchange (IKE) — The first phase of IPSec authentication, which accomplishes key management. IKE is a service that runs on UDP port 500. Once IKE has established the rules for the type of keys two nodes will use, IPSec invokes its second phase, encryption.

Internet Protocol Security (IPSec) — A Layer 3 protocol that defines encryption, authentication, and key management for TCP/IP transmissions. IPSec is an enhancement to IPv4 and native to IPv6. IPSec is unique among authentication methods in that it adds security information to the header of all IP packets.

IP spoofing — A security attack in which an outsider obtains internal IP addresses, then uses those addresses to pretend that he or she has authority to access a private network from the Internet.

Kerberos — A cross-platform authentication protocol that uses key encryption to verify the identity of clients and to securely exchange information once a client logs onto a system. It is an example of a private key encryption service.

key — A series of characters that is combined with a block of data during that data's encryption. In order to decrypt the resulting data, the recipient must also possess the key.

key distribution center (KDC) — In Kerberos terminology, the server that runs the authentication service and the ticket granting service in order to issue keys and tickets to clients. On a Windows 2000 network, a user's domain controller serves as his or her KDC.

key management — The method whereby two nodes using key encryption agree on common parameters for the keys they will use in order to encrypt data.

key pair — The combination of a public and private key used to decipher data that has been encrypted using public key encryption.

Layer 2 Forwarding (L2F) — A Layer 2 protocol similar to PPTP that provides tunneling for other protocols and can work with the authentication methods used by PPP. L2F was developed by Cisco Systems and requires special hardware on the

host system end. It can encapsulate protocols to fit more than just the IP format, unlike PPTP.

Layer 2 Tunneling Protocol (L2TP) — A Layer 2 tunneling protocol developed by a number of industry consortia. L2TP is an enhanced version of L2F. Like L2F, it supports multiple protocols; unlike L2F, it does not require costly hardware upgrades to implement. L2TP is optimized to work with the next generation of IP (IPv6) and IPSec (the Layer 3 IP encryption protocol).

packet-filtering firewall — A router that operates at the Data Link and Transport layers of the OSI Model, examining the header of every packet of data that it receives to determine whether that type of packet is authorized to continue to its destination. Packet-filtering firewalls are also called screening firewalls.

Point-to-Point Tunneling Protocol (PPTP) — A Layer 2 protocol developed by Microsoft that encapsulates PPP so that any type of data can traverse the Internet masked as pure IP transmissions. PPTP supports the encryption, authentication, and LAN access services provided by RAS. Instead of users having to dial directly into an access server, they can dial into their ISP using PPTP and gain access to their corporate LAN over the Internet.

Pretty Good Privacy (PGP) — A key-based encryption system for e-mail that uses a two-step verification process.

principal — In Kerberos terminology, a user.

private key encryption — A type of key encryption in which the sender and receiver have private keys, which only they know. Data encryption standard (DES), which was developed by IBM in the 1970s, is a popular example of a private key encryption technique. Private key encryption is also known as symmetric encryption.

proxy server — A network host that runs a proxy service. Proxy servers may also be called gateways.

proxy service — A software application on a network host that acts as an intermediary between the external and internal networks, screening all incoming and outgoing traffic and providing one address to the outside world, instead of revealing the addresses of internal LAN devices.

public key encryption — A form of key encryption in which data are encrypted using two keys: one is a key known only to a user, and the other is a key associated with the user and can be obtained from a public source, such as a public key server. Some examples of public key algorithms include RSA (named after its creators, Rivest, Shamir, and Adleman), Diffie-Hellman, and Elliptic-curve cryptography. Public key encryption is also known as asymmetric encryption.

public-key server — A publicly available host (such as an Internet host) that provides free access to a list of users' public keys (for use in public key encryption).

Remote Authentication Dial-In User Service (RADIUS) — A server that offers authentication services to the network's access server (which may run the Windows NT or 2000 RAS or Novell's NAS, for example). RADIUS provides a single, centralized point of authentication for dial-in users and is often used by ISPs.

root — A highly privileged user ID that has all rights to create, delete, modify, move, read, write, or execute files on a system. This term may specifically refer to the administrator on a UNIX-based network.

screening firewall — See *packet-filtering firewall*.

Secure Sockets Layer (SSL) — A method of encrypting TCP/IP transmissions—including Web pages and data entered into Web forms—en route between the client and server using public key encryption technology.

security audit — An assessment of an organization's security vulnerabilities. A security audit should be performed at least annually and preferably quarterly or sooner if the network has undergone significant changes. For each risk found, it should rate the severity of a potential breach, as well as its likelihood.

server_hello — In the context of SSL encryption, a message issued from the server to the client that confirms the information the server received in the client_hello message and agrees to certain terms of encryption based on the options the client supplied. Depending on the Web server's preferred encryption method, the server may choose to issue your browser a public key or a digital certificate at this time.

session key — In the context of Kerberos authentication, a key issued to both the client and service by the authentication service that uniquely identifies their session.

social engineering — Manipulating relationships to circumvent network security measures and gain access to a system.

SSL session — In the context of SSL encryption, an association between the client and server that is defined by an agreement on a specific set of encryption techniques. An SSL session allows the client and server to continue to exchange data securely as long as the client is still connected to the server. SSL sessions are established by the SSL handshake protocol.

symmetric encryption — A method of encryption that requires the same key to encode the data as is used to decode the cipher text.

Terminal Access Controller Access Control System (TACACS) — A centralized authentication system for remote access servers that is similar to RADIUS.

ticket — In Kerberos terminology, a temporary set of credentials that a client uses to prove that its identity has been validated by the authentication service.

ticket granting service (TGS) — In Kerberos terminology, an application that runs on the key distribution center that issues ticket granting tickets to clients so that they need not request a new ticket for each new service they want to access.

ticket granting ticket (TGT) — In Kerberos terminology, a ticket that enables a user to be accepted as a validated principal by multiple services.

Transport Layer Security (TLS) — A version of SSL being standardized by the Internet Engineering Task Force (IETF). With TLS, IETF aims to create a version of SSL that will encrypt UDP as well as TCP transmissions. TLS, which will likely be supported by new Web browsers, uses slightly different encryption algorithms than SSL, but otherwise is very similar to the most recent version of SSL.

tunneling — The process of encapsulating one protocol to make it appear as another type of protocol.

REVIEW QUESTIONS

1. If you have root privileges on a system, you could delete user IDs from that system. True or False?
2. What do you call manipulating people to get them to reveal confidential information, such as their passwords?
   a. social engineering
   b. social manipulation
   c. social coercion
   d. social affectation
   e. social posturing
3. Which of the following is the most secure password?
   a. 123
   b. dolphins
   c. !tz0g557x
   d. tchotchke
   e. 1040506
4. Which two of the following would not typically be used for authenticating via a network operating system?
   a. IP address
   b. user name
   c. password
   d. last name
   e. date of last logon
5. Name three different security risks associated with people.
6. What is the most likely way that a network's security will be compromised?
   a. from within the organization
   b. from a cracker on the Internet
   c. from a cracker posing as a contractor
   d. from a cracker using IP spoofing over a modem connection
   e. from a cracker using a remote control program
7. Which device could a cracker use to intercept and interpret transmissions between one router and another router on a WAN?
   a. router
   b. hub

c. switch

d. sniffer

e. multimeter

8. Accepting the default options for security on a server-based application is usually a good policy. True or False?

9. If someone obtains one of your LAN's internal IP addresses and uses it to gain access through your firewall from the Internet, what method of security attack is he or she using?

   a. flashing

   b. SSL

   c. denial of service

   d. framing

   e. IP spoofing

10. The UDP protocol is more secure than the TCP protocol. True or False?

11. If someone floods your LAN's router with excessive traffic so that your legitimate traffic cannot go out or come in, what method of security attack is he or she using?

    a. flashing

    b. SSL

    c. denial-of-service

    d. framing

    e. IP spoofing

12. Which of the following is not typically addressed in a security policy?

    a. preventing accidental damage to network software and hardware

    b. specifying what model and make of firewall is appropriate for the network

    c. ensuring that authorized users have appropriate access to the resources they need

    d. communicating each employee's responsibilities with respect to maintaining data integrity

    e. identifying a response to suspected security breaches

13. What is the primary purpose for establishing a security response team?

    a. to demonstrate to users the security risks they face

    b. to train users to respond to security threats as they happen

    c. to devise a coordinated response to security breaches while or after they occur

    d. to comprehensively audit the security of the network

    e. to publicize the sanctions that will befall users who do not follow the security policy

14. What should an organization do to assess its potential security risks?
    a. perform a security audit
    b. freeze any new user ID creation
    c. question users about odd behavior on their workstations
    d. train users to watch for suspicious activity
    e. hire a number of different consultants to provide multiple perspectives on the best approach to network security
15. Name four questions that should be addressed in a security audit.
16. What's the simplest way to stop a denial-of-service attack on a server?
    a. shut down the victimized server
    b. shut down the firewall between your server and the Internet
    c. turn on TCP and UDP filtering
    d. restart your central switch to clear traffic from the server's segment
    e. install a sniffer on the server to filter out traffic issued from the denial-of-service attacker
17. Which of the following transmission media is the most secure?
    a. UTP
    b. STP
    c. coaxial cable
    d. infrared
    e. fiber-optic cable
18. Which of the following encryption methods is most commonly used on VPNs?
    a. SSL
    b. Kerberos
    c. IPSec
    d. PGP
    e. private key encryption
19. Which two of the following do not contribute to a network's physical security?
    a. closed-circuit TV
    b. public keys
    c. badge access systems
    d. door locks
    e. digital certificates

20. Which of the following network operating system restrictions is most likely to stop a cracker who is attempting to discover someone's password?
    a. number of unsuccessful logon attempts
    b. time of day
    c. total time logged on
    d. source address
    e. username and password
21. Name four criteria that a packet-filtering firewall might use for filtering traffic.
22. At which two layers of the OSI Model does a packet-filtering firewall operate?
    a. Transport and Network layers
    b. Network and Data Link layers
    c. Data Link and Transport layers
    d. Session and Transport layers
    e. Physical and Data Link layers
23. Before a firewall can effectively filter unwanted traffic, it must be:
    a. placed between a private and public network
    b. configured according to an organization's security needs
    c. combined with a proxy server
    d. attached to a switch on the internal LAN
    e. installed with Kerberos server software
24. Which of the following best describes the function of a proxy server?
    a. to deny LAN access to specific IP addresses
    b. to filter inappropriate content traveling from the Internet to an internal LAN
    c. to encapsulate protocols in the IP format
    d. to act as a gateway between an internal LAN and the outside world, masking the IP addresses of private LAN devices
    e. to issue and retain public keys for users requiring PGP e-mail security
25. Which of the following security risks does using the callback feature on a remote control application address?
    a. the possibility that passersby can take over the host system that is being controlled remotely
    b. the possibility that unauthorized users can take over the host system after they have discovered its phone number and logon ID
    c. the possibility that unauthorized users can scan a host system's ports and discover its phone number

d. the possibility that passersby can shut down the system in mid-session

e. the possibility that a cracker can gain access to the system through an intermediate host, such as the file server

26. If a company wants to save office leasing costs and allow 50 of its employees to work at home, what type of arrangement would be the most secure, practical, and economical for granting home workers access to the LAN?

a. create exceptions in the firewall filtering rules to accept incoming traffic from each home worker's workstation according to their IP addresses

b. set up a VPN that uses a RADIUS server to centrally authenticate and grant LAN access to dial-in users

c. establish a Windows 2000 RAS server with direct PPP dial-in capability

d. program Web front ends for the applications used by home workers and employ SSL to transmit their work over the Internet

27. What service does PPTP provide?

a. It encapsulates protocols so they can run on IP-based networks such as the Internet.

b. It authenticates dial-up users according to their source address.

c. It ensures that remote control callback mechanisms are secure.

d. It encrypts data using the IPv6 protocol.

e. It tracks suspicious IP activity on a packet-filtering firewall.

28. If you are entering your account number in a Web page to gain access to your stock portfolio online, which of the following encryption methods are you most likely using?

a. PGP

b. Kerberos

c. L2F

d. IPSec

e. SSL

29. In general, the longer the key, the more secure the encryption. True or False?

30. PGP is frequently used for what type of network communication?

a. e-mail

b. FTP

c. HTTP

d. Telnet

e. HTTPS

HANDS-ON PROJECTS

Project 15-1

For a networking professional, it's important to stay abreast of new security threats and learn how to address them. In fact, in a large organization, a team of professionals might be devoted to network security, with one team member responsible for researching new security threats. In this project, you will look at some Web resources that can help you find out about vulnerabilities on your network. For this project, you will need a workstation with Internet connectivity and a Web browser.

1. Connect to the Internet and point your browser to the following URL: **www.microsoft.com/security/bulletins/current.asp**. The Security Bulletin Search page appears. Scroll down the page until you find the list of security risks associated with Microsoft software according to the date they were discovered.
2. Scroll through the list and click **MS01-024, "Malformed Request to Domain Controller can Cause Memory Exhaustion."**
3. Read the description of the problem and how Microsoft has addressed it. How was this problem discovered and reported to Microsoft? How could Windows 2000 Server allow someone to exploit this vulnerability?
4. Click the **Back** button on your browser and browse more Microsoft security bulletins.
5. Now point your browser to this URL: **developer.novell.com/research/appnotes/2000/june/03/a0006037.htm**. Read about Novell's recommendations for NDS security.
6. According to what you read in the NetWare security document, how does the concept of "inherited rights" affect users within an organizational unit? Under what circumstances does this document suggest that using network address restrictions on a user account would be helpful? Why does this document recommend that the network administrator have at least two separate accounts: one with Administrator privileges and one with normal user privileges?
7. One organization that provides an updated list of many types of security risks is CERT, a clearinghouse for security risks established by the Carnegie Mellon Software Engineering Institute. To view its current alerts, point your browser to the following URL: **www.cert.org/advisories/**. Notice that the alerts are organized by the date they were released.
8. View information about a denial of service alert by clicking one of the most recent bulletins.
9. Read the advisory. What type of action does CERT recommend network administrators take to defend against or prevent this threat?
10. Click the **Back** button on your browser to return to the list of advisories. Browse through the most recent alerts. To what types of software or systems do most of the alerts pertain?

Project 15-2

As you have learned, password restrictions play a significant role in network security. In Chapter 8 you learned how to impose password restrictions on a user account as you created it. In this project, you will learn how to impose security restrictions on all user accounts within a domain at once. For this exercise, you will need a Windows 2000 server and the capability to log on as Administrator to that server. To test your changes, you will need a Windows 2000 Professional or Windows 9x client and a valid user ID (other than the Administrator) that can log onto the Windows 2000 server.

1. Log onto the server as Administrator.
2. Click **Start**, point to **Programs**, point to **Administrative Tools**, then click **Domain Security Policy**. The Domain Security Policy window appears.
3. In the left-hand pane of the Domain Security Policy window, double-click **Security Settings**, if necessary, to expand the container.
4. Double-click **Account Policies**, if necessary, to expand this option.
5. Click the **Password Policy** option. A list of password policies and their settings appears in the right-hand pane of the Domain Security Policy window, as shown in Figure 15-9.

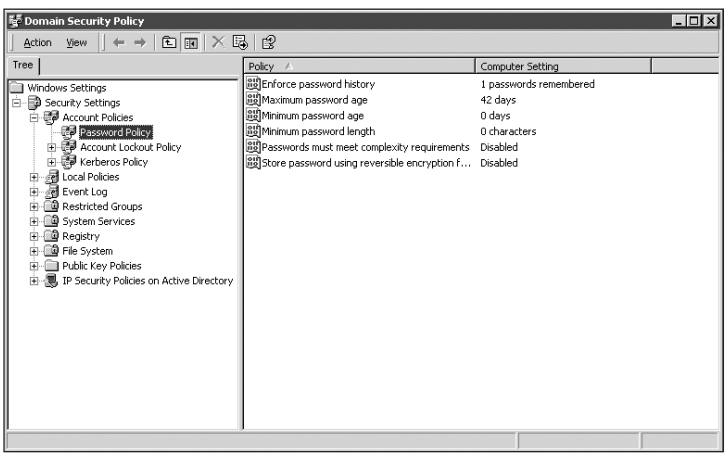

Figure 15-9 Windows 2000 Domain Security Policy window

6. What is the maximum password age set at? Right-click on the **Maximum password age** policy, then click **Security** from the shortcut menu. The Security Policy Setting dialog box appears.
7. Change the number of days after which passwords will expire to 60 days.
8. Click **OK** to close the Security Policy Setting dialog box.

9. What is the minimum password length setting? Right-click the **Minimum password length** policy, then click **Security** in the shortcut menu. The Security Policy Setting dialog box appears.
10. Change the minimum number of characters in the password to 8.
11. Click **OK** to close the Security Policy Setting dialog box.
12. Right-click on the **Passwords must meet complexity requirements** policy, then click **Security** in the shortcut menu. The Security Policy Setting dialog box appears.
13. Make sure the Define this policy setting option is checked, then click the **Enabled** radio button to enable the complexity requirements. Click **OK** to close the Security Policy Setting dialog box.
14. In the left-hand pane of the Domain Security Policy window, click the **Account Lockout Policy** option. Account lockout policies appear in the right-hand pane of the Domain Security Policy window.
15. Right-click the **Account lockout duration** policy, then click **Security** in the shortcut menu. The Security Policy Setting dialog box appears.
16. Check the box next to **Define this policy setting**, then change the number of minutes an account will be locked out to 5.
17. Click **OK** to close the Security Policy Setting dialog box. A Suggested Value Changes window may appear, informing you that because the account lockout duration is now 5 minutes, the settings for the "Account lockout threshold" and "Reset account lockout counter" will also be changed. Why do you think the number of invalid logon attempts was raised when you lowered the lockout duration?
18. Click **OK** to accept the changes and close the Suggested Value Changes window.
19. Close the Domain Security Policy window.
20. From your Windows workstation, attempt to log onto the Windows 2000 server with an ordinary user ID. Are you prompted with any messages about your password?
21. Now attempt to change your password, using each of the following character strings: dog, 12345, and TJ01xxN73. How does the server respond to each?
22. Log off from the server, then attempt to log on again, but deliberately enter the wrong password five times in a row. Then try the correct password on your sixth logon attempt. What happens?
23. Wait six minutes and try logging onto the server with the correct password. What happens?

Project 15-3

Another important principle of protecting network data from security breaches is assigning the proper rights to each individual or group that has access to your network's servers. In this project, you will assign appropriate rights for five groups of users on a Windows 2000 server. You will need a Windows 2000 server with a Windows 2000 Professional client workstation attached and the capability to log onto the server. You should have Administrator rights on the server. The Windows 2000 server should contain the following user IDs: Bob, Patrick, Mary, Errol, Sally, Chris, Inez, Richard, Dave, and Cory. Each ID should be associated with the same password, "3netch15309", and users should have no modifications to their default file access rights. The server should also contain the following directories:

C:\DATA\BUDGET

C:\DATA\SAMPLES

C:\DATA\CONTRACTS

C:\DATA\PAYABLES

C:\DATA\RECEIVABLES

The users should belong to the groups specified in Table 15-1:

Table 15-1 Users and groups for Project 15-3

Users	Group
Bob, Patrick, Mary, Errol, Sally, Chris, Inez, Richard, Dave, Cory	Accounting
Mary, Patrick, Sally	Accounts Payable
Bob, Chris, Cory	Accounts Receivable
Inez, Richard, Chris	Finance
Errol, Patrick, Dave, Inez	Managers

1. Log onto the server from the client workstation as Patrick, using the password 3netch15309.
2. Attempt to open the directory C:\DATA\CONTRACTS on the server. What message do you see?
3. Now you will give group rights to each directory. Begin by logging onto the Windows 2000 server as Administrator.
4. Double-click the **My Computer** icon to see a list of drives on the server.
5. Double-click the **Local Disk (C:)** drive icon to view its contents.
6. Double-click the **DATA** directory to view its subdirectories.
7. Right-click the **BUDGET** folder, then click **Properties**. The BUDGET Properties dialog box opens.
8. Click the **Security** tab.

9. Click **Add**. The Select Users, Computers, or Groups dialog box opens.
10. Double-click the **Accounting** group, then click **OK**. You return to the BUDGET Properties dialog box.
11. With the Accounting group highlighted, check the **Write** box in the Allow column. This setting gives the Accounting group permission to create (or add) new files, but not to modify existing files.
12. Click **OK** in the BUDGET Properties dialog box.
13. Repeat Steps 7 through 12 using the directory names and their respective privileges as shown in Table 15-2:

Table 15-2 Directories and group permissions for Project 15-3

Directory	Group	Permissions
C:\DATA\SAMPLES	All groups	Modify
C:\DATA\CONTRACTS	Managers	Full control
C:\DATA\CONTRACTS	Accounts Receivable	Read (only)
C:\DATA\CONTRACTS	Accounts Payable	Read (only)
C:\DATA\PAYABLES	Accounts Payable	Modify
C:\DATA\RECEIVABLES	Accounts Receivable	Modify

CASE PROJECTS

1. As an experienced networking professional, you have been asked to conduct a security audit on a local credit union's network. The union currently has two locations, a headquarters office downtown and a branch office on the east side of town. The headquarters has the following equipment:
   - 20 Windows 2000 Professional workstations, connected to a Windows 2000 server
   - 1 Windows 2000 RAS server accessed by home workers after hours
   - 1 Windows 2000 server for recordkeeping
   - 1 UNIX database server
   - 1 UNIX Web server for members to check their account balances online
   - 1 firewall where the network connects to the credit union's ISP via a T1 dedicated link

   The east-side office has five Windows 2000 Professional workstations, connected to the headquarters office Windows 2000 server through a dedicated ISDN link.

All tape backups are housed in a secure room in the headquarters office, with copies being kept in a file cabinet at the east-side office. At the headquarters, the servers reside in a locked room that admits authorized users with an electronic badge access system. Both locations have numerous security cameras, including cameras in the computer room and backup tape storage vault at the headquarters. The manager also tells you that the credit union has a security policy that all employees are required to read and sign when they become employees. He believes that the network is very secure and asks you if he could do anything else to ensure that the network is safe from security breaches. In response, create a checklist of items on this network that should be evaluated for security. Describe any access points or situations that constitute potential security risks. In addition, explain how the credit union manager could better train his employees to understand network security.

2. As part of your security audit, you have recommended that credit union employees change their passwords so they are more secure and that the IT department enforce password changes every 60 days, because of the confidential nature of the data on the workstations and servers. The credit union employees are not enthusiastic about this change, and they complain that they already have too many things to remember. How might you convince them that choosing secure passwords and changing their passwords frequently are in their own best interest and for the good of their employer?
3. The credit union is experiencing tremendous growth and needs to either open another branch office on the west side of town or allow their auditors and loan-processing staff to work from home. It asks you to compare the security requirements of opening a new branch office versus implementing a dial-VPN solution (using the Internet) for work-at-home employees. As part of your comparison, list the costs associated with these security requirements. For an expansion of 10 users, which solution do you recommend?